Kr00k

Description: This week we look at a significant milestone for Let's Encrypt; the uncertain future of Facebook, Google, Twitter and others in Pakistan; some revealing information about the facial image scraping and recognition company Clearview AI; the Swiss government's reaction to the Crypto AG revelations; a "must patch now" emergency for Apache Tomcat servers; a revisit of OCSP stapling; a tried and true means of increasing your immunity to viruses; an update on SpinRite; and the latest serious vulnerability in our WiFi infrastructure, known as Kr00k.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-756.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-756-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Lots to talk about. We've got Let's Encrypt's one-billionth certificate, just in time for a major security problem that they're rushing to fix. We'll have the details on that. Also, the new Kr00k vulnerability at WPA3. Steve has a prescription that'll keep you safe. And we're going to talk about Clearview AI, the face recognition that was used by a lot more people than anybody thought. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 756, recorded Tuesday, March 3rd, 2020: Kr00k.

It's time for Security Now!, the show where we cover your security, your privacy, your health, your happiness, your welfare, online and off. In fact, this show is the one show you need to get everything you need for the rest of your life. That's Steve Gibson right there. He's the man in charge, giving you the Vulcan salute. He's at GRC.com. Note we both have our "I Voted" stickers.

Steve Gibson: Yes, we do.

Leo: Nice job. Nice job. I got up early because I had to go before iOS Today. So I got up early. It's a nice feeling. You go, and it's in a school. The volunteers are all there. I said, "Thank you for your service and helping us vote Super Tuesday."

Steve: And down here every single person in the area received a mail-in ballot. Was that California-wide, or only in Orange County?

Leo: No, it might be California-wide. I mean, I've always been an absentee voter.

Steve: Yeah, as have I.

Leo: But I got my ballot, and I had the good sense to wait. I marked it, which is a good thing because the candidate I voted for had dropped out by the time the election happened. What happens then? They still get the votes. And then I guess they can give their delegates...

Steve: And so they're able to sort of like say, oh, wow, we would have done better than we thought. But, you know...

Leo: Too late.

Steve: Yeah.

Leo: Too late.

Steve: So we have Episode 756, which I want to get over with, Leo, so that I can go watch the election returns for the rest of the day.

Leo: Watch TV, yes.

Steve: But we're not going to hurry this along. We're going to talk about Kr00k, K-R-0-0-K, for March 3rd, 2020. Kr00k is a recently discovered bad vulnerability in the WiFi WPA2 protocol.

Leo: Oh, no. So I had a caller on Sunday that said, "Oh, this WiFi, I don't want to use it." I said, "Oh, KRACK, don't worry about it. It's hard to do. The guy's got to be sitting out front." Now there's Kr00k.

Steve: There's now Kr00k, which it's an extension of KRACK. And it affects more than a billion devices. But first we're going to look at a significant milestone for Let's Encrypt; the uncertain future of Facebook, Google, Twitter, and others in Pakistan; some revealing information about the facial image scraping and recognition company we've spoken of several times, Clearview AI; the Swiss government's reaction to the Crypto AG revelations; a must-patch-now emergency for Apache Tomcat servers; a revisit of OCSP stapling; a tried-and-true means of increasing, of all things, your immunity to viruses. I don't mean computer viruses, I mean human viruses.

Leo: Oh, good. We need that.

Steve: It was apropos of the coronavirus problem that the world is facing right now. A quick update on SpinRite and my ongoing development of it, and then we're going to look at this latest serious vulnerability in our WiFi infrastructure known as Kr00k. I also have the most bizarre Picture of the Week...

Leo: Yes, you do.

Steve: ...we've ever had. And a pointer to our listeners to the coolest kinetic sculpture on Kickstarter I've ever seen.

Leo: You're going to have to be careful on this show because apparently I must be hormonal or something. But I want to buy everything. And I even want to buy your Picture of the Week. I want to buy that.

Steve: It's called "Being Leo."

Leo: No, I don't always feel this way. Oh, man.

Steve: So we always have a Picture of the Week. And normally it involves, typically, security somehow.

Leo: Uh-huh.

Steve: Someone sent this to me, and it's just so bizarre that I thought, well, okay.

Leo: I want it.

Steve: Yeah. So I think what it must be is a very creative right-angle plug.

Leo: Yeah, it's an elbow, yeah.

Steve: When you want - exactly. When you want a cord to be plugged into the wall, but not to be sticking straight out, you want it at a 90-degree angle. And so for those who don't have the benefit, or maybe we're sparing them, video of what we're seeing right now, this is a nose which has been plugged into the wall, and then the cord is shoved up its nostrils.

Leo: I'm buying these. I'm putting these everywhere in the house. This is awesome.

Steve: It is. I mean, and think about people who come over will do a double-take. It's like, okay, well, okay, consider where we are. We're at Leo's. So, yeah, we're going to have nose plugs everywhere. So anyway, it's wacky. It's apropos of nothing. But it does look like it's a real thing. It looks like you could go, I mean, I didn't even look on Amazon to see if I could find a right-angle nose plug for my wall socket.

Leo: If I find it, I will - because I am absolutely looking for it. If I find it, I'll let you know. It's hysterical.

Steve: Okay. But on a more serious note, also apropos of nothing, because I am, you know, I sort of keep tabs on what is going on in Kickstarter, I was informed this morning of the coolest - not inexpensive, but for the right person, this is the right thing. I created one of my GRC shortcuts to help people find it: grc.sc/sand, S-A-N-D. So this is a 12-inch round - it's also available in a star shape - 12-inch round shallow tray filled with fine-grain white sand. There is a steel ball bearing rolling around in the sand under control of a microprocessor and an arm system underneath such that the ball can be made to roll wherever it wants. It leaves a trail. And through a succession of movements over time, this builds up just fabulous patterns in the sand.

So again, it's apropos of nothing except I love our listeners, and this thing will go away. If it's going to sell out, I'd rather it sold out to people listening to this podcast than other random people who are less deserving.

Leo: I want this so badly.

Steve: So to our deserving listeners - oh. And, now, Leo, you've got the mechanism on screen right now, which I didn't even...

Leo: You bought it before you even got to the bottom of the page.

Steve: I so immediately knew I had to have one of these that I didn't even scroll down to see the mechanism underneath, which is wonderful, in order to move the little ball bearing around. So for what it's worth, for our listeners, grc.sc/sand will redirect you to this Kickstarter campaign that looks like the real deal. All of the required caveats about, yeah, well, it's not commercially made, who knows if it'll ever happen, blah blah blah. Okay, fine.

Leo: Steve.

Steve: Don't buy it if that's a deal breaker for you.

Leo: You should build this yourself. You could write the software to make the arm...

Steve: Well, actually, Leo, when I retire - because Lorrie will go bonkers when she sees this. I got this for her. We're going to make a full-size rectangular coffee table that has ball bearings rolling around in it, doing this.

Leo: Honestly, this would be a fun project. A Raspberry Pi, you've got to get the arm somehow, but...

Steve: It'd be perfect. And, you know, because I'm never going to give up assembly language. And so this will be something I can do in my, you know, in retirement. But in the meantime, we're going to have one because - oh, and there is, by the way, an add-on. You can't buy it right now. They've already done twice what they were hoping to achieve. But they are limited in the quantity that they're producing. So for what it's

worth, grc.sc/sand. I hope you're listening to this in time, if it's the kind of thing you would like to have.

Leo: $349 for the birch one, which is not bad. Dark walnut, $399. Not bad.

Steve: I mean, it's a work of art.

Leo: It's beautiful.

Steve: It's a beautiful thing.

Leo: He's not making much money on this because the amount of time, just the woodworking alone...

Steve: Yeah.

Leo: That's a lot of work. Wow. That is really nice. Good tip.

Steve: Just a tip for our listeners. So Let's Encrypt last week hit a milestone. They issued their billionth, with a "b," billionth certificate. And what's interesting is the timescale of this. I thought this was an interesting opportunity to remind our listeners of just how much reluctance and inertia this industry has. It's just, I mean, we're always talking about it. We're seeing instances of it all the time.

Netscape was the originator of SSL, the Secure Socket Layer, which was an addition; it provided authentication and privacy thanks to encryption, over the existing TCP connection. So your browser, your Netscape Navigator, would connect to a server with a standard TCP connection. And then once they had connected, the browser and the server would negotiate an encryption layer on top of that existing TCP connection to create an authenticated and private tunnel through which they could then communicate. That happened, and that added the "S" to the HTTP, giving us HTTPS in - wait for it - 1994 is when that happened. But its adoption is a repeat of this classic tale of Internet adoption reluctance that we keep seeing. I have in the show notes a table I found showing the percentage...

Leo: Wait a minute. I pushed the wrong button. I did find the nose wall outlet for $35. It's from This Is Why I'm Broke. It's kind of the story of my life, Steve Gibson. All right. I'll get the other one.

Steve: It does exist. So it's a real thing.

Leo: It's a real thing.

Steve: A nose outlet. Okay. Anyway, so I found a table - okay, now. So HTTPS became possible in 1994. Let's see. What would it be? It would be 19 years, no, 21 years later, in

August of 2015 - 21 years later, August of 2015, HTTPS was at 6.71%. Half a year later, 9.39. August of 2016, 17.76%. Six months later, February 2017, nearly 20%. Okay, in 2017, February of 2017. We just left February of 2020. So three years ago, three years ago it was one in five connections were secure, 19.96%. That's how recent this shift has been. So, and in fact it was when the connections crossed the 50-50 point there was some notice of that. That was the summer of 2018. Only as recent as that. Which is just amazing to me, again, to look back and see.

And this is the Alexa top million sites that these stats were generated from. So of course there's no doubt that it was a combination of factors that got together to finally, after what did I say, 21 years of, like, nobody really caring? And in fact, remember when we were talking about Firesheep. Sites had the ability to switch you into HTTPS when they wanted to solicit a username and password. But then, inadvisably, they would switch you back to HTTP, even though the cookie that you had obtained to create a persistent session was now in the clear. If you were HTTP, everybody could see the session cookie, and that's where Firesheep demonstrated how easy it was to just commandeer somebody's session. In any situation where you were open WiFi, like I keep using Starbucks as an example, you could just say, oh, and just take over their session by looking at their cookie and using it yourself. And you would now be logged in as them.

So, and of course it was Snowden, the Snowden revelations of how prevalent the actual eavesdropping of things that were going on was, that finally caused the industry to get off its butt and take privacy, Internet connection and communications privacy more seriously. I mean, there's...

Leo: And a little credit to Google because didn't they say we're not going to - if a site's not HTTPS, we're going to not give you as high a ranking in the search rankings?

Steve: Yes, that was...

Leo: I think that must have been the reason; right?

Steve: I have that in my show notes.

Leo: Oh, good. Okay, sorry.

Steve: So, oh, no, no, it's good. For one thing, Let's Encrypt solved the...

Leo: That made it easy; right.

Steve: Yeah, well, it made it free, which was one of the arguments. It was like, I mean, there were old-school - you and I were talking about how old we are. But we're not crusty, really, yet. There were crusty Unix people that just, as a matter of principle they said, "I'm not paying for a certificate." You know, it's like...

Leo: [Cranky sounds]

Steve: The Stallman types, it's like [cranky sounds]. You know, it's like, no. So, okay. Now you don't have to pay anything, thanks to Let's Encrypt. And then of course there's the annual hassle argument. It's like, well, if I just could buy it once, that'd be fine. I mean, I was in that camp once upon a time. It's like [cranky sounds]. I mean, like, I'm just - I'm paying for bits. You know, I'm not paying for anything. Except now I really do appreciate the fact that I am paying DigiCert for them taking the time to verify I am who I am claiming I am for the certificate that I use.

So Let's Encrypt, as we know, was started as a nonprofit effort by the ISRG, the Internet Security Research Group, which was multiply sponsored by the EFF, the Electronic Frontier Foundation; Cisco, Facebook, Google, the Internet Society, those are the IETF guys; Mozilla; and, strangely enough, a French cloud service provider, OVH. They're the backers behind Let's Encrypt. And it's been, by every measure, after a billion certificates issued, an overwhelming success.

This gave us ACME, which is the reverse-engineered acronym, had to be, Automated Certificate Management Environment is what ACME stands for. ACME is the automation protocol which verifies your domain ownership and then automates the issuing of a certificate to a bot that you've got on your server that accepts the certificate and installs it. And so this all now happens without you having to do anything. As we know, Let's Encrypt's certificates have a 90-day lifetime because, thanks to freeness and automation, once you set it up, it just runs by itself, and it takes care of everything.

So the other impetus is, as we've been covering on the podcast, our browsers are beginning to incrementally shame and strengthen their shaming of any and all non-HTTPS web pages, even going so far as to say that, if you don't have HTTPS, the page is not secure, which technically is true, but there are pages on the 'Net you could argue that have absolutely no need for security. But doesn't matter. Moving forward, browsers will be shaming websites.

And of course, as you noted, Leo, Google went one step further, and they added whether a page was delivered over a secure connection as one of many signals which they funnel into their page-ranking algorithm because, I mean, you could argue that a site that is delivering all of its pages over a secure connection may be more worthwhile. I mean, it's arbitrary, but Google's trying to find things to pull together in order to create a ranking. So it's one of the things that Google has officially said they're using to rank pages. So sites that want a stronger rating will be using a certificate of one form or another.

So anyway, I know that I may seem to be harping on the issue of Let's Encrypt and the fact that, sort of to remind people that they are automating the proof of domain validation only. And they are issuing certificates for any domain, including, for example, Playpal.com. And I guess since no one but we geek techies have ever really understood the distinction among DV, OV, and EV certs, I do see the logic of zeroing the assertion being made by domain validation certs, meaning that if you have a bot which is issuing certs to anybody who can prove that they control a domain, then that cert by definition no longer means anything other than it's providing security, it's providing encryption for that domain. Yes, it is verifying, you know, it's providing authentication to the domain. But the domain could be anything. I mean, it could be a forged spoofing site, but at least now it's a secure forgery. So, good.

I did a little bit of research because I was sort of curious. Research conducted back on March 20th of 2017 - so here we are March of 2020, three years ago - revealed that Let's Encrypt had among their billion certificates that they've issued, issued 15,270 PayPal certificates, meaning certificates either containing the term PayPal or some visual lookalike phrase, obviously being used to spoof the authentic PayPal domain. PayPal's certificate, I went over and looked at it last night when I was putting this together, is an extended validation EV certificate, obtained, as is mine, from DigiCert. So imagine how

the real PayPal feels about having 15,270 lookalike certificates issued to secure spoofing domains created only to confuse their users.

And this is why I believe it would be useful and meaningful to users to have our web browsers eventually indicate when a website is protected. Yes, it's protected. It's encrypted. But it's a certificate that was obtained by an organization that some human took a few minutes to verify, like PayPal or like GRC. Again, nothing against Let's Encrypt except the fact that they're automated, 100% automated, unfortunately means that they are a target for spoofers. And obviously not just it could happen, but it is happening.

We've talked a couple times about Russia, and laughed at the small size of the threats that the government is making to our social media, our Western social media companies when they're disobeying Roskomnadzor, the Russian authority that's going to say, okay, we want you to move all your servers into Russia to serve Russian citizens. And, you know, there have been some skirmishes. But, you know, if it's however many millions of rubles, it ends up being pocket change for any of these larger social media companies.

Well, another instance of this, not Russian, but this case the government of Pakistan has published some proposed stringent censorship rules governing online content that have Facebook, Google, and Twitter collectively threatening to pull up stakes and go entirely dark throughout the country of Pakistan.

Leo: What? Oh, my god.

Steve: Yeah. Yeah. This is known, the new regulations are known as Pakistan's "Citizens Protection Against Online Harm Rules for 2020." I've got a link to the actual rules in our notes for this episode. The rules laid out by the government of Pakistan give authorities the power to demand social media platforms remove any content they deem questionable within 24 hours. And to that end, Pakistan has proposed the creation of a National Coordinator Office to monitor the content of online services. Additionally, social media platforms must provide a way to prevent the live streaming of "online content related to terrorism, extremism, hate speech, defamation, fake news, incitement to violence, and national security."

Leo: Well, that doesn't sound too bad. I mean, that's...

Steve: Well, except they actually have to provide a means to do it. And so within three months of the new rules coming into play, companies such as Facebook and Twitter must also open up permanent offices in the country, establish one or more local servers to store data in Pakistan, and must also agree to "remove, suspend, or disable access to such account, online content of citizens of Pakistan residing outside its territorial boundaries and posts on online content that are involved in spreading of fake news or defamation and violates or affects the religious, cultural, ethnic, or national security sensitivities of Pakistan." I mean, that's what it says in the rules.

The rules also give the government the right to block a social network if they refuse to comply, or impose fines of up to 500 million rupees, which in this case is more than $50. It's approximately 6.9 million USD. So it's a penalty that has a bit of teeth.

Anyway, those fighting for free speech online argue that such wide-reaching powers are designed to curb free speech and impose censorship. And notice that this is just - the way this is written, it's that requests that are made must be honored. So they're not - Pakistan is not saying, "We want you to use your best judgment." They're saying, "We're

going to give you a list, and within 24 hours you must remove anything that we tell you we don't like."

Facebook, Twitter, and Google are agreeing that this is curbing free speech and imposing censorship. Those three organizations are part of something known as the Asia Internet Coalition, the AIC, which is a trade association discussing issues surrounding Internet innovation and regulation in the region. In response to Pakistan's proposed rules, this AIC group have replied: "The rules as currently written would make it extremely difficult for AIC Members to make their services available to Pakistani users and businesses." In other words, and there has been some explicit conversation to this effect, upwards of 70 million Pakistani residents could find themselves unable to access Facebook and Twitter, and denied the use of any of Google's wide range of services, which are in wide use by businesses now around the world.

The AIC response went on to add that: "AIC members recognize Pakistan's strong potential, but the sudden announcement of these rules belies the government of Pakistan's claims that it is open for business and investment. As no other country has announced such a sweeping set of rules, Pakistan risks becoming a global outlier, needlessly isolating and depriving Pakistani users and businesses of the growth potential of the Internet economy." The AIC added that it wished to work with the government of Pakistan to come up with more appropriate solutions to online data and content management without risk of crippling of Pakistan's emerging digital economy.

The New York Times, which first reported on this, noted that by threatening to leave altogether, the companies may be attempting to apply pressure - you think? - to Pakistan's government to quickly rethink the proposed rules or face protests from the country's citizens and business owners when the services are withdrawn from the country. Yeah. However, Pakistan is not alone. Last year India also proposed a set of rules in the same vein, prompting the same concerns over free speech. India is expected to publish their guidelines soon.

So we have a bit of a mess on our global hands. Our Western social media companies were created in a comparatively permissive and open environment of the United States, where we enjoy a great deal of freedom. But even here in the U.S. we're seeing increasing trouble with online hate speech, counterfactual content posted and posing as factual, and the fact that "news" is increasingly available from unvetted sources. And of course we have the continuing encryption problem. So this is just, as I keep saying, Leo, a really interesting time for...

Leo: It's challenging because we have notions of free speech that are unique to our country. We have the First Amendment.

Steve: Yes. Yes.

Leo: And of course we think this is all a really good idea. And we maybe deal with the consequences. I mean, there's free speech on Twitter. So I don't - it seems like the height of arrogance to tell a country, well, no.

Steve: To impose our standards.

Leo: This is what we think, and you've got to do it this way.

Steve: Yup.

Leo: So it's challenging because at the same time we do think free speech is a good thing. But I can see how a country might not, especially in the face of terrorism and fake news and hate speech and all of those things. We've got that stuff in spades on Twitter. And there's no...

Steve: And, you know, I sort of hail from a "technically how do you pull it off" perspective. I mean, so you have a country like Pakistan that says, essentially is telling the purveyors of these services, "We demand full right to censor the content you provide in our country." That's what they're saying.

Leo: But how do you do it; right. How do you put it into implementation?

Steve: Yes, exactly, exactly. So you set up in some office in the Pakistani government a portal that allows anything they wish to be blanked from delivery within Pakistan's borders. Okay. Then, I mean, I guess that's what you do. Or you decide we're not going to do that, and you withdraw your service. Well, it seems unlikely because these services are generating huge revenues for our Western social media companies. I would imagine that what's going to evolve is exactly that, that the governments are given full censorship rights of any content that they choose to censor within their borders to their people. And I would imagine in time Facebook and Google and Twitter will kick and scream, but they will end up capitulating and be willing to be there, whatever content survives.

And, you know, if that happens, and there is strong censorship of content, then people who want to have their content not blocked are going to have to behave themselves, or a very powerful censor is going to descend on them with both feet. So it'll be interesting to see how this happens. I mean, like how it goes. I really do wonder whether these, I mean - and maybe there's a compromise. But I don't see how. I mean, Google doesn't want the responsibility of making those decisions themselves. We watched Facebook struggling with this whole issue of what to censor and what not. I mean, recently in the news has been, like, oh, yeah, just, what was it, pour Clorox over yourself, and you won't get the coronavirus.

Leo: Yeah, right. They've got to ban that.

Steve: Yeah. It's not a good idea.

Leo: No.

Steve: To allow fake health news from happening.

Leo: It's a tough one.

Steve: Really, really interesting.

Leo: Really. And we knew this. This is the culture clash that we knew a global Internet would offer. And, you know, the answer is not to just say, well, everybody has to be like us. Right? But I don't know what the solution is. And especially, you're right, the technological solution is very challenging. But remember we said that about the right to be forgotten. How hard is that going to be to implement; right?

Steve: Yeah. So our "friends" at Clearview AI, Leo. Remember that they're - and I have "friends" in quotes. They're the company, just for clarity's sake, they're the company who was scraping facial content from the web and reselling it to unnamed but presumably highly interested third parties, without the knowledge or permission of those whose faces have been captured by their database. Well, in a somewhat, well, interesting turn of events, the unnamed purchasers are unnamed no longer because last Wednesday Clearview revealed that it was the victim of a data breach through which it lost control of its customer list.

Leo: Okay. Right? Okay.

Steve: Whoopsie.

Leo: Whoopsie. They didn't get - the database of faces did not leak out, but that's just chance; right?

Steve: So along with the information, including the number of searches those customers have made and how many accounts those customers had set up. So remember that the app identifies people by comparing photos to a database of images scraped from social media and other sources. It first rose to our attention earlier this year when a New York Times investigation into the software company revealed its activities. At that time, the Times revealed that Clearview AI had quietly sold access to faceprints and facial recognition software to more than 600 law enforcement agencies across the U.S., claiming that it could identify a person based on a single photo and turn over their real name and far more information about them.

Within a few weeks of the Times article, Clearview was being sued by a potential class action lawsuit that claims the company amassed the photos out of "pure greed," although I really don't know what difference their motivation makes, to sell to law enforcement, thereby violating the nation's strictest biometrics privacy law, which as we know is in Illinois, the Biometric Information Privacy Act, BIPA, where any biometric data, including a person's own face, cannot be captured and used without their explicit permission. And of course, as we said at the time, Senator Edward Markey, learning of this, called Clearview AI a "chilling privacy risk."

Since then, Facebook, Google, YouTube, Microsoft, and Twitter have all sent cease-and-desist letters to Clearview AI, saying that we do not give you permission to scrape our content, our public content, for your purposes. But as we talked about when we talked about this last time, Leo, you brought up the correct fact that the question of automated scraping has been moving around throughout the courts, and the most recent appellate decision came down on the side of the scraper, saying that, hey, if it's publicly presented content, it's available for any purpose, including scraping.

So the question now, who exactly has been dipping into the Clearview well? Well, not only the expected law enforcement agencies like ICE, our Immigration and Customs Enforcement agency in the U.S., and the U.S. Department of Justice, people at the FBI,

Customs and Border Protection, and Interpol, but also, according to reporting from BuzzFeed News that obtained access to this list, AT&T, Verizon, T-Mobile, Best Buy, Eventbrite, Las Vegas Sands...

Leo: Eventbrite?

Steve: I know. Eventbrite. That's the one that I said, what?

Leo: I can see why a casino or a store, maybe. But Eventbrite?

Steve: I know. Coinbase, Bank of America, Walmart, Kohl's, and Macy's.

Leo: Oh, boy.

Steve: The privacy geeks are all naturally freaked out. Nathan Freed Wessler, a staff attorney with the American Civil Liberties Union, our ACLU, said: "This list, if confirmed, is a privacy, security, and civil liberties nightmare. Government agents should not be running our faces against a shadily assembled database of billions of our photos in secret, with no safeguards against abuse." For its part, Interpol confirmed that a small number of its officers in the Crimes Against Children Unit used 30-day trials of the Clearview AI product, but that Interpol has no formal relationship with the company. In an email statement, ICE confirmed its use of Clearview AI, saying it's primarily for agents with Homeland Security Investigations who are involved in child exploitation and cybercrime cases. The FBI declined to comment. And note that, Leo, as usual, everyone is hiding behind the children.

Leo: Yeah, yeah, of course.

Steve: Saying that, oh, well...

Leo: We only did it to save the kids.

Steve: That's right. AT&T said it's not a client of Clearview AI. Best Buy denied ever using or planning to use Clearview AI. BofA said it's not a customer. Eventbrite denied being a client of the company. Those not responding to requests for comment were the Department of Justice, Customs and Border Protection, Verizon, T-Mobile, Las Vegas Sands, Walmart, Kohl's, and Macy's.

Oh, but that leaves Coinbase. They said they hadn't made a commitment to use Clearview AI. In an email statement a spokesperson said: "Our security and compliance teams tested Clearview AI to see if the service could meaningfully bolster our efforts to protect employees and offices against physical threats and investigate fraud."

Leo: Oh. That might make sense, yeah.

Steve: They said: "We've not tested nor would we use Clearview AI's service with our customer data. We maintain strict privacy controls that prevent customer data from being used in this manner." BuzzFeed's report also said Clearview AI has expanded to law enforcement agencies in, and I sorted these into alphabetical order: Australia, Belgium, Brazil, Canada, Denmark, Finland, France, Ireland, India, Italy, Latvia, Lithuania, Malta, the Netherlands, Norway, Portugal, Serbia, Slovenia, Spain, Sweden, Switzerland and the United Kingdom. So seems to be a pretty popular service. Business appears to be brisk at Clearview AI. And facial recognition through social media content scraping and analysis looks like it's giving us a brave new world indeed.

Leo: I feel like, because they offer such long free trials, that a lot of people would try it; right? I want to try it.

Steve: Well, and you know, if for example Coinbase saying that they're wanting to - presumably they've got cameras.

Leo: In their office, yeah.

Steve: In their offices.

Leo: Yeah, that makes sense.

Steve: But that would suggest that it actually flags you as a criminal because instead of, like, your place of employment and your name, I mean, what they want is immediate go/no-go, do you have - does this person have a police record.

Leo: No. I'll give you an example. So the way I understand Clearview AI works is you would give it a picture, and it would say that's who this is. It does not - I don't think it gives you, oh, yeah, it's a felon. It just says these are 15 other examples of that person. We believe this is their personal information. Imagine you have a disgruntled employee, and you have a camera out front, and you say, hey, let me know if this guy shows up. That's an example of use.

So I think there's a lot of offices that that wouldn't be an unusual use of it, especially in this day and age with disgruntled employees coming back with firearms. That might not be an unusual use of it. And again, if you're in the hacked database because you did a free trial, I mean, I could have been in that. I mean, who wouldn't be in it; right?

Steve: Yeah. So this suggests that they are making a lot of money. And that argues...

Leo: Oh, yeah, I'm sure they are.

Steve: ...that they're going to be able to defend themselves legally. And, you know, we're in an unsettled area of the law right now, like whether or not, you know, this is going to prove to be - and, you know, as you and I have talked about, just sort of capitulating to this whole problem.

Leo: What are you going to do? Yeah.

Steve: Yeah, well, you know, I do have a face, and I like to wear it in public. So I'm going to be captured on camera. And when Lorrie and I did our SQRL European tour, I created an Instagram account, and there we were, smiling from the Eiffel Tower. So, whoops. I'm sure I'm part of the Clearview database now.

Leo: Yup. Yup, you are. Mm-hmm. Back to you, Steverino, now fully caffeinated with about a gallon of fine Starbucks brew.

Steve: Recaffeinated, yes.

Leo: Good lord. Now, that's a Pringle's tin with a Starbucks wrapper around it. Tell the truth, Steve.

Steve: Actually, in order to keep it warm, it's several of their venti paper cups stacked.

Leo: Okay, but still...

Steve: Oh, yeah, well, it's...

Leo: It's still pretty tall. It's just a venti? It feels like it's bigger than a venti.

Steve: A quinti venti, yes.

Leo: Quinti venti, there you go.

Steve: Yeah. There we go. So not surprisingly, I think you brought it to our attention a couple weeks ago because the news was just breaking as we were going to the air, the Crypto AG organization, which was a Swiss-based supplier of cryptographic technology, had been just at that time discovered to have been owned by the U.S. CIA and the equivalent German organization.

Leo: Just a little thing. Just a little thing.

Steve: And it's like, oh, my goodness. So it's predictably blown up with apparently well-founded allegations that U.S. and German intelligence deliberately implemented backdoors in Crypto AG's systems for the purpose of eavesdropping on governments worldwide.

Leo: No.

Steve: Governments being the customers of Crypto AG. Switzerland has filed a complaint. The project was known as Operation Rubicon. And as a result of this investigation by the Washington Post, ZDF, which is a German public service television broadcaster, and SRF News, which is a Swiss radio station, those are the three entities that pooled their resources and discovered what had been going on with this Swiss company, Crypto AG. The report that has been filed says that governments were paying "good money to the U.S. and West Germany for the privilege of having their most secret communications read by at least two, and possibly as many as five or six, foreign countries."

This further suggests that communication records may have been shared between the Five Eyes members that we've spoken of, that's the U.S., the U.K., Canada, Australia, and New Zealand. So that five plus Germany, since Germany was also a party to ownership of Crypto AG. Being a Swiss company, Crypto AG covert operations, the fact of this has shaken Switzerland, due to its long-term standing as a neutral country in political matters, a reputation that it has defended avidly for years.

In 2018, Crypto AG split into two companies: CyOne AG, C-Y-O-N-E, CyOne AG, which serves the local Swiss market; and Crypto International AG. After the investigation was published, the Swiss Ministry of Economic Affairs prohibited exports of Crypto AG products, basically shutting it down. CyOne AG, that only works within Switzerland, has maintained that it is independent of its now ill-reputed international counterpart.

The Swiss government has appointed a former Supreme Court judge to investigate the matter. His report is expected this summer, in June. The Swiss attorney general's office said on Sunday that the criminal complaint recorded against "persons unknown" has been formally filed by the State Secretariat for Economic Affairs, which is SECO in Switzerland.

Reuters reported that the Swiss attorney general's office will review the complaint and ultimately decide whether or not to open a criminal investigation. The investigation may be focused upon determining who within the company knew about the surveillance practices in the hopes of mitigating damage caused by Switzerland's neutral position. You know, they're going to say, hey, this wasn't the whole company, this wasn't everybody, it was just a select few people who had knowledge of this.

It's interesting, too, because there has been rumor for quite a while that there was a way for the U.S. and Germany - it's just sort of been in the air, never confirmed, no clear facts to back it up. But just sort of some off, you know, some occasional rumors that maybe this crypto wasn't as strong as was believed. And of course, unfortunately, when the report comes out with details that back up what has sort of been believed, but nobody was quite sure about it, then that puts that issue to rest.

So anyway, it'll be interesting to see how this goes. And I don't know who will step in to fill the vacuum. Maybe, again, this company has been around for decades. So it's probably just been inertia which has allowed it to continue long after there were alternative good sources of crypto. They acquired, being Swiss, they had a reputation for honesty and integrity, which unfortunately does not look like it was deserved.

There is a web server technology that has been around for a long time, 22 years in fact. Tomcat is an open source implementation. Tomcat runs under Apache, the Apache server. It's an open source implementation of the Java Servlet, Java Server Pages, Java Expression Language, and Java WebSocket technologies. So it provides a pure Java HTTP web server environment for hosting Java-based web applications. It is very popular, especially in the enterprise where Java is often the chosen implementation language. It started off 22 years ago, in 1998, as a reference implementation of a Java servlet which was created by James Duncan Davidson, who was a software architect at the time at Sun

Microsystems. Since then, it has been evolving steadily across a series of minor and major releases. It is a real deal.

To give our listeners a sense for its penetration, Shodan lists more than 890,000 Tomcat servers currently reachable over the Internet - 890,000 - and the similar BinaryEdge service has located more than a million. So right now there are more than a million Tomcat servers vulnerable to a newly discovered problem, an exploit known as "Ghostcat." Ghostcat is a high-risk file read-and-include vulnerability being tracked as CVE-2020-1938. It's present in all the Apache JServ Protocol (AJP) of Apache Tomcat between versions 6.x and 9.x, which is all versions of Apache's Tomcat server released during the past 13 years. So 13 years, a million-plus servers. And this AJP protocol is exposed publicly by default when Tomcat is brought up. It is a critical flaw that can lead to a server takeover.

The Tomcat developers, who are doubtless embarrassed by this discovery, in a somewhat vain attempt to put the best possible face on this disaster, wrote: "Tomcat treats AJP connections as having higher trust than, for example, a similar HTTP connection. If such connections are available to an attacker, they can be exploited in ways that may be surprising." Yeah, surprising to the people whose servers are taken over remotely by hackers on the Internet.

Researchers at the Chinese security firm Chaitin Tech, who discovered the bug, explained that after successfully exploiting an unpatched Tomcat server - which again, any Tomcat server using Tomcat in the last 13 years. They said: "An attacker can read the contents of configuration files and source code files of all web apps deployed on Tomcat. In addition, if the website app allows users to upload files" - so, okay, that's good. So it's maybe not all, but many sites provide some means, some reason for allowing uploads - "an attacker can first upload a file containing malicious JSP script code to the server."

They said: "The uploaded file itself can be any type of file, such as pictures, plaintext, et cetera." In other words, it could be a JPEG. It could be any kind of a Java-based web app that accepts customer-based posts, which, for example, a social media platform that allows you to upload your own avatar for the account that you create, for example. So it doesn't have to - it can be masquerading as a JPEG, and then this thing is then able, as soon as you get the code onto the server, that is, you get a file onto the server, the exploit allows that to be treated as JavaScript, or JSP script code, which allows it then to be exploited using the Ghostcat vulnerability, turning it into a remote code execution vulnerability.

According to Snyk and Red Hat, Tomcat also ships with apps built using the Spring Boot Java framework, as well as other Java-based servers and frameworks including but not limited to the JBoss Web Server (JWS) and JBoss Enterprise Application Platform (EAP). And not surprisingly, given the power of this exploit, coupled with the fact that the enterprise targets lying behind these Java-based servers are likely quite attractive.

And I have a link here because I was curious who's using this. There is a cwiki.apache.org for the Tomcat Apache add-on. They have a "powered by" where they're bragging about which organizations are powered by Tomcat. And it's a little hair-raising because there are many significant organizations that number themselves among Tomcat users.

Leo: It used to be a big deal. I'm surprised this many people still use it. I felt like...

Steve: Inertia, Leo. Inertia.

Leo: Inertia. Once Bergen Jersey Foreclosures is working, you don't want to change it. [Crosstalk] by Tomcat for years.

Steve: Yes, and I noted [crosstalk].

Leo: It says.

Steve: Yup, yup. Also, shoot, there was a stock trading company.

Leo: eTrade. It says eTrade on there.

Steve: eTrade, I know, eTrade is there. It's like, whoopsie.

Leo: But it may be, oh, I should just click this because it may be they used to - well, it doesn't bring it to the article. Maybe they used to use it.

Steve: So, and the threat is not just theoretical. The cyberthreat intelligence firm Bad Packets tweeted on Saturday: "Mass scanning activity targeting this vulnerability has begun. Patch now." So if any of our listeners have their organizations using Tomcat on Apache, yeah. I mean, and certainly you're using something from the last 13 years. Update yourself. All the versions, the latest 6.x, 7.x, or maybe it's 7.x, 8.x, and 9.x. 6.x is no longer in service. So it's 7.anything, 8.anything, and 9.anything. And they have updates for all three branches. So whichever one you're using, definitely update.

I wanted to revisit briefly our discussion of OCSP Must-Staple that we talked about last week, Leo. You and I came to agreement on-air that, yeah, it was such an obvious solution to the problem of certificate revocation. So I was thinking about that some more. I hadn't looked at it in a long time. But I was wondering why it might not be the right solution. Was there some downside that I wasn't aware of?

Leo: Right.

Steve: And so it occurred to me that if it was up to the web server to instruct the web browser to require a freshly stapled certificate, then a stolen and later revoked certificate might still be usable because the malicious server would certainly not ask browsers to require that the stolen certificate be accompanied with a fresh stapled OCSP assertion. So I double-checked. I found the well-known security researcher Scott Helme's clearly written page on OCSP Must-Staple. I have a link in the show notes for anyone who's interested. His page starts out by saying: "Revocation checking is broken and has been for some time." I think this was a couple years ago, as I recall. He said: "While some vendors have sort of worked around this with proprietary solutions, there is little that smaller sites can do.

"In the early days of the web we had Certificate Revocation Lists, or CRLs. These were lists of all certificates that a CA had revoked and could be downloaded by a client to check if the certificate they were served had been revoked. These lists didn't scale and eventually downloading these large files became a problem. Thus the Online Certificate Status Protocol, or OCSP, was born. Instead of the client downloading a list of all revoked

certificates, they would submit a request to the CA to check the status of the specific certificate they have just received. Sadly, OCSP was riddled with problems like poor CA infrastructure being unavailable and the privacy concern of clients leaking the site they were visiting back to the CA." That's a good point. I had forgotten to mention that last week.

"To get around this problem, OCSP Stapling was created. Instead of the client making the OCSP request to the CA, the host website would make the request and 'staple' the response to the certificate when they served it. Because the OCSP response is short-lived and digitally signed by the CA, the client can trust the stapled OCSP response. The final problem was that the client had no idea that the site in question supports OCSP and whether or not it should expect them to staple an OCSP response." That last is what had occurred to me. He says: "Thus we finally arrived at OCSP Must-Staple." That is, not just OCSP stapling, but OCSP Must-Staple.

"Setting the OCSP Must-Staple is fairly easy as it's simply a flag that needs to be set by your certificate authority in the certificate they generate for you. This flag instructs the browser that the certificate must be served with a valid OCSP response, or your browser should hard fail the connection." So in other words, that's exactly what we wanted. It's not up to the server. The must-staple flag is bound into the certificate authority-signed certificate. So any attempted use of the certificate would carry the must-staple flag. As long as the recipient honors the must-staple flag, the certificate would be far better protected than even a one-year expiration on certificates would provide.

So the only reason then I could see for not taking the must-staple path would be that we're not yet ready for its deployment. I did some more digging. And I found a year-old assessment of the support of OCSP Must-Staple among browsers. As of January 2019, a little over a year ago, only Firefox 60 of all browsers respected OCSP Must-Staple.

Leo: So what's the point?

Steve: Yes. Chrome 66 did not. Opera did not. Safari does not. IE11 did not. Edge 42 did not. And the last update on the server side that I was able to find indicated that, while Microsoft's IIS supports OCSP Stapling properly, neither Apache nor Nginx do. Those open source server implementations were still minimal and not well implemented. So our industry has taken the path of least resistance and has chosen the lowest common denominator solution, which doesn't even begin to deliver equivalent security. One year? We could easily have one day of window. But because no one has pushed the open source servers nor end-users' browsers, with the sole exception of Mozilla's Firefox, everyone is instead having to be hassled with much more frequent certificate issuance, whether manually or through automation. And even after that's done, we still have truly pathetic certificate revocation security. A year, as we noted last week.

And you know, you really would think that Google, with their proven ability to guide the industry due to the strength of Chrome and now Chromium, which everybody else is using except Mozilla, you'd think that Google would be leading the way here. After all, their browser's revocation with that CRLSET debacle remains, as we know, the worst of all. But that hasn't happened yet.

So it really looks as though this is where we as an industry should focus. Once our web servers are up to speed, and all our browsers honor Must-Staple, then nothing - and think about this. Once our web browsers are up to speed with honoring Must-Staple, which could easily be done, and we have web servers which - again, IIS is there, and Apache, and Nginx. They sort of have half-hearted implementations. They just need to

get fixed. They don't do the - they don't look at the Staple expiring and go out and reach out and respect the cache life of the certificate and so forth. So trivially fixed.

If we had that, then nothing would prevent certificate authorities from having a choice, from giving users a choice. They could issue 10-year life certs with must-staple flags set, which would be far more secure than anything we have today. And browser behavior could be adaptive. Apple could set up Safari to limit the total lifetime of all non-must-staple certificates to one year. And maybe it even can get shorter in the future. But then to allow any certificate bearing the must-staple flag to have any lifetime it wishes because the second it becomes revoked, within a day it would stop being honored. Nothing is preventing that scenario other than the industry's own laziness and inertia.

Leo: You know who wishes it were working right now? Let's Encrypt. They have to revoke three million certificates. And they gave users 24 hours to fix it.

Steve: Yup. You need to refresh your Let's Encrypt cert. They found a bug in their CAA code.

Leo: Whoops.

Steve: Which is the DNS record that allows a Certificate Authority permission to issue a certificate. We've talked about this before, the CAA record. It specifies which certificate authorities are permitted to issue that domain certificates. So it is a voluntary thing; but it's simple, and simple to implement. And again, Let's Encrypt had a bug in their code. So you're right, Leo.

Leo: So if they revoke these certs today, which they're going to do, they say, but we don't have a system for checking revocation, does it matter?

Steve: Yeah. The sites will stop working.

Leo: Oh, they will stop working.

Steve: Oh, I see. If they revoke them, yes, and nobody checks.

Leo: And nobody checks.

Steve: I wonder if Let's Encrypt runs an OCSP server. They probably do.

Leo: Oh, maybe they staple and say check our server. Maybe they do. Because otherwise, I mean, isn't that the problem, that you can revoke all you want? It doesn't much matter?

Steve: Yeah. Yeah.

Leo: Since nobody's checking? No, they must do something else. I wonder if they do, if they run their own...

Steve: The good news is anyone using Chrome will have no problem at all.

Leo: Because it will just go [dismissive sound].

Steve: Because Chrome does not check. It doesn't care at all.

Leo: Doesn't care at all.

Steve: No.

Leo: Oh, my god. It could not be worse.

Steve: So I have a weird public service announcement, which I was put in mind of by all this recent news about the COVID-19 coronavirus pandemic. Which put me in mind of Security Now! Episode 200 - I was going to say 2,000, but no, it'll only have three digits.

Leo: No.

Steve: 209, which you and I recorded in mid-August of 2009.

Leo: Have it been that long? Wow.

Steve: So 10.5 years ago, Leo.

Leo: That's amazing. It feels like just yesterday.

Steve: And I went to the episode. It describes itself, it says: "Steve and Leo kick off the podcast's fifth year with a rare off-topic discussion of something Steve has been researching for the past eight weeks and passionately believes everyone needs to know about." And that of course was the famous Vitamin D episode.

Leo: Okay. So last time we talked you said, oh, you don't have to take as much as we thought.

Steve: I don't remember saying that.

Leo: Oh, okay. Okay. I must have misunderstood.

Steve: I don't remember saying that. So, you know, I don't want to go on at great length because all the information is available. GRC.com/health/vitamin-d.htm will quickly take you to the audio. We also covered news of the week and so forth. But I finished by talking about Vitamin D. You can also just google "GRC Vitamin D." It'll take you to that page. Any listeners who joined us in the last 10.5 years probably missed that. And so that's why I'm bringing it up today, and only just pointing to it. The short version is I explained in great detail what I had learned, the fact that it's not a vitamin. It's been misnamed, mislabeled. It's actually a hormone. It's fat soluble, not water soluble, so it builds up over time in our tissues.

The government's stated minimal - the RDA, the Recommended Daily Allowance, I think last time I looked was 400 IU, which is pathetic. It doesn't - it's not - it's just ridiculous. And we have no way of really knowing what we should have, except that by looking at the blood of people who work with their shirts off in the sun, literally lifeguards and roofers, we believe we have a sense for what a natural blood level is. The problem is it's synthesized. There's almost no dietary source. Our skin synthesizes it when UVB radiation interacts with cholesterol in our skin to produce this hormone, to synthesize it as a result of UVB radiation hitting our skin. But UVB radiation doesn't penetrate glass. We are typically behind glass or with clothes on most of the time.

Leo: Well, nobody goes outside anymore. If you do, you slather sunscreen on, and you put on a broad-brimmed hat.

Steve: You see people hiding behind umbrellas to keep the sun off of their skin.

Leo: My alabaster skin may not be touched by the sun.

Steve: So what happened was I actually got an interesting research study out of this because we talked about it in mid-August 2009. In the spring of 2010 I began getting email from our listeners who were incredulous over the fact that it was the first winter they had ever in their memory gone through where they never got sick. People, their family members were getting sick around them. They didn't get sick because they listened to the podcast. They heeded the advice. I love this advice because a bottle of 5,000 IU of Vitamin D, 360 little tiny, little itty-bitty, I call them "little drops of sunshine," costs about $15. So a year's supply, one a day, a year's supply is 15 bucks.

And I'm bringing it up now because it is an across-the-board immune booster. The podcast goes into lots of detail about how Caucasians happened as a result of the importance of Vitamin D. So I just wanted to sort of point our listeners to it. We've got this problem right now with what looks like a viral pandemic that's going to be significant. So this is when you want you and your family and your friends to have the greatest immunity to viruses possible. One thing you can easily do is think about Vitamin D. Listen to the podcast.

Leo: 5,000 IU.

Steve: 5,000 IU.

Leo: Is there a toxic level? I mean, there's something we shouldn't go...

Steve: Yes, yes. If you were to take 50,000 IU per day for several months, you could reach toxic levels.

Leo: So only take one, not 10.

Steve: Just take one.

Leo: Okay.

Steve: Yes. Now, and for example, physicians who encounter people who are deficient will give somebody a shot in the arm of 25,000 IU. So a brief burst of it also won't hurt you.

Leo: It ain't gonna kill you.

Steve: It literally takes 50,000 a day for months for you to hurt yourself.

Leo: But still, let's not overdo it.

Steve: And you should - there's no need to take more than 5,000 IU. When I was experimenting, I was at 10. And after a few months I noted my blood level because I was testing every week, I noticed that my level in my blood of the active form of Vitamin D, 25OHD, would finally kind of be creeping up into the 80s or 90s. We don't even really know for sure what a maximum problem is. But 5,000 is all you need. 400, which is what the government recommends, is not getting you off the ground. Oh, and there's also D2 versus D3. What you want - I explain all this in the podcast. So if I've tickled your interest, listeners, and you're thinking about the coronavirus, and you're not already taking 5,000 IU of Vitamin D a day, check it out.

Leo: D2, not D3.

Steve: No, no, D3.

Leo: D3, not D2.

Steve: Yeah, D2 is not really effective. It's synthesized from lanolin, as I recall; and it's not effective, whereas D3 is. So anyway, don't take my word for it. Do some googling. In the years since I keep getting, like, there will be studies that are coming out, and they all reaffirm that we're not getting the Vitamin D we should. A $15 bottle for a year just makes so much sense. And for what it's worth, I nor my friends never get sick anymore because after I learned this, I switched to 5,000 IU a day, and it just ended. You know, I tended to have a strong immune system anyway. But now more than ever it really makes sense to beef that up.

Leo: Or 125 micrograms, if you live in the sensible world.

Steve: Yes, yeah. And that's just it. It is, I remember talking about it, you need so little bit of it that...

Leo: 125 micrograms is nothing.

Steve: Yes. They use a huge vat of olive oil. They, like, drop a dropper into this huge vat. They make sure that they stir it up sufficiently, and then they make these little itty-bitty capsules, so even people who are pill phobic, who can't swallow pills, these things are just little teeny bitty drops.

Leo: Oh, they're teeny, yeah.

Steve: Yeah.

Leo: I just put them in with all the other horse pills I'm taking. I never even notice.

Steve: Me, too. So I wanted to mention that I started working on the AHCI driver for SpinRite. I already have low-level BIOS-bypassing hardware drivers in place for all the pre-AHCI technologies, so traditional IDE and the Legacy SATA controllers. That's done now. So the final piece of technology, as I mentioned last week, will be to add the same low-level hardware support for AHCI controllers.

We won't initially have USB, NVMe, or UEFI support. Those will follow later. But giving SpinRite the highest possible speed support for all spinning and SATA drives of any size will fulfill my commitment to bring the aging SpinRite v6 current with today's latest technologies. And that's got my attention. So I just wanted to give our listeners an update and thank them for their patience.

Leo: How exciting. Do you think Vitamin C will have any help in that, as well?

Steve: Well, I didn't want to push it because taking a useful amount of Vitamin C is difficult.

Leo: It's a lot. But I do what Lorrie does. I put the liquid Vitamin C in my water.

Steve: Yes. Yes, that liquid Vitamin C is what I would recommend for people. The problem with C is that it's water soluble. So it doesn't stay in us. On the one hand, you don't have to worry about overdosing because you just - your body will just eliminate it.

Leo: That's why you have to take it, instead of taking a big pill, you want to sip a little bit of it in the water at a time.

Steve: And that's a great way to do it. And in fact that C that Lorrie and you are taking...

Leo: Delicious.

Steve: ...has a nice sort of a citrus flavor to it.

Leo: [Crosstalk] orange water. It's really good, yeah. I'm very happy.

Steve: Yeah. And that's what I would recommend, yup.

Leo: I get three grams a day in my 54-ounce Bubba.

Steve: Good. And if you were taking 20 to 25, then you'd be taking enough.

Leo: Steve says I have subclinical scurvy.

Steve: That's true.

Leo: How much?

Steve: Yeah, about 20,000 grams. I mean, 20,000 milligrams, 20 grams.

Leo: 20 grams a day.

Steve: Yeah, that's about what your liver is trying to produce. It's a six-stage process to convert glucose into Vitamin C. The last stage, there's an enzyme, L-gulonolactone oxidase, which we are not synthesizing, and that keeps our liver from giving us Vitamin C. It's trying to. Dogs and cats all make at least 10 grams a day.

Leo: Wow.

Steve: So anyway.

Leo: You should have as much as a dog, anyway.

Steve: That's right.

Leo: All right. I'm going to put more - I'm going to have more of that stuff.

Steve: Do it. Especially now.

Leo: It does give you chapped lips. I've got to tell you, it gives you chapped lips.

Steve: I didn't know that.

Leo: Did you not? Have you not experienced that?

Steve: No, because I just take tablets. I'm a tablet-taker.

Leo: Oh, I'm a tablet-taker.

Steve: But if you only do one thing, do Vitamin D. If you want to do more, then see about the liquid form of Vitamin C, which makes it very easy to sip throughout the day.

Leo: Yeah, yeah. Kr00k is KRACK.

Steve: So every so often a vulnerability comes along, and I'm not happy about it. I'm not celebrating it. But it is just so, well, our listeners will find out in a second. It's just perfect. Researchers at ESET named their latest discovery Kr00k, K-R-0-0-K, to highlight the zeroes in the name because the zeroes is what this latest discovery of a very serious WiFi vulnerability affecting more than a billion devices is all about. So I really shouldn't sound happy about this at all. And I'm not. I mean, this is bad. It's not, you know, end-of-the-world bad. But I'll explain it in a minute.

It affects anything not yet patched containing the most popular WiFi chips in the industry, which are those manufactured by Broadcom and Cypress. And those chips are included in devices such as Amazon's consumer Echo and Kindle devices; Apple's iPhones, iPads, and MacBooks; Google's Nexus smartphones; Samsung's Galaxy smartphones; the Raspberry Pi 3, Xiaomi Redmi smartphones, and access points from Asus and Huawei. I mean, just to give a sense for how widespread this is. All of those are, or were, vulnerable to Kr00k. In other words, all consumer devices, access points, routers, and IoT gadgets, anything using WiFi chips from Broadcom or Cypress. So we know, just based on the numbers, that this weighs in at more than a billion WiFi-capable things. And since the Cypress chips are popular among IoT devices, their numbers are untold.

Okay. So what happened? ESET researchers came up with a very clever hack, which I suppose I'd call an "interlock" hack because it leverages two things that should have been carefully interlocked, but weren't. A communications buffer that should have been flushed, but wasn't being, can arrange to be flushed if the device's firmware is updated. So that's what will change after the firmware update.

But the bad news is that any of these more than a billion devices that are not updated will forever remain vulnerable to this subtle but readily exploitable firmware design flaw. And you know, it seems like we're talking about these kinds of problems more and more. You remember we had a Bluetooth, a similar Bluetooth problem not long ago, and I was using the example of an alarm system that was Bluetooth-enabled, and it had been abandoned by its Chinese supplier, or they'd gone out of business, or who knows what.

Or there was just never any intent to ever update its firmware after it was sold. So it would forever be vulnerable.

So as our listeners are doubtless aware, this idea that everything is vulnerable at some level has gradually become an important and recurring topic of this podcast. And as in this instance, this belief is not based upon a fear of the unknown. It's unfortunately based upon solid practical experience and evidence. Given everything we've learned about the desires and motivations of hackers and state actors, there's just no question that somewhere, or more likely at many somewheres, there are people methodically compiling these increasing number of things that, for some large swath of the device population, are never going to be fixed. You know, for some time now I've been characterizing security as being porous, but it's becoming more so than I ever expected.

Leo: Poor us.

Steve: Yes, poor us.

Leo: Poor us.

Steve: Yes. And so like all of the worst and most powerful attacks, this one is not complicated. It's not like Spectre and Meltdown and those sort of theoretical...

Leo: KRACK wasn't that easy, either; was it? I mean, it was...

Steve: Nah, it wasn't. And this is even easier, Leo. Get a load of this. Okay. So point-to-point 802.11 Ethernet WiFi operates by associating and disassociating the paired endpoints with each other. The association process involves the establishment of a shared cryptographic key given a shared secret. So, for example, our home router is preconfigured with our WiFi network password. So we give the same password to our laptop computer. And because each knows the same shared secret, they're able to negotiate a session key which will be used to encrypt the communications.

We've talked about WiFi attacks in the past. Being radio, denial of service is easy. One brute force approach would be to jam the receiving end with radio frequency noise. But a much more subtle attack is made possible by the fact that 802.11 WiFi connection management packets are never encrypted. By definition, they are in the clear. So they're like the management, not the actual connection. So it's possible to deny any unwitting WiFi user the use of the nearby WiFi access point, simply by spoofing and sending disassociate packets to their laptop. When the laptop receives a disassociate command, which again is not encrypted, so there's no need to know what the network's key is, that command will be immediately obeyed, and the user will drop off that 'Net connection. However, their endpoint will then immediately start working to reassociate itself with the access point, so this process would need to be repeated if you wanted to hold them in denial of service.

Okay. The clever Kr00k vulnerability occurs as a side effect of that disassociation. Inside the WiFi firmware is the negotiated encryption key, and a fairly sizable 32K byte transmission buffer. As one would hope and want, when a WiFi session is disassociated, that transient session encryption key that they were using is immediately and proactively zeroed. It is wiped and literally set to all zeroes. But unfortunately that 32K buffer is not also flushed. Instead, believe it or not, it continues being sent out of the WiFi radio

transmitter under the now zero key. And the way the WPA2 AES-CCMP cipher mode operates, that's the standard cipher mode for WPA2, an all-zero key results in no encryption. So the 32K bytes of data that was supposed to be protected under the shared session key now emerges from the device's WiFi radio transmitter in the clear as plaintext, completely unprotected.

ESET's disclosure whitepaper said: "This serious flaw, assigned CVE-2019-15126, causes vulnerable devices to use an all-zero encryption key to encrypt part of the user's communication. In a successful attack, this vulnerability allows an adversary to decrypt some wireless network packets transmitted by a vulnerable device. Kr00k affects devices with WiFi chips," they write, "by Broadcom and Cypress that haven't yet been patched. These are the most common WiFi chips used in contemporary WiFi-capable devices such as smartphones, tablets, laptops, and IoT gadgets. Not only client devices, but also WiFi access points and routers with Broadcom chips were affected by the vulnerability, thus making many environments with unaffected or already patched client devices" - like our smartphones - "vulnerable anyway."

In other words, yeah, if you have an Apple iPhone, Apple has already patched. But what about your old router? And what about the router at Starbucks, or at the airport, or on the plane? They said, continuing with what they wrote: "Our tests confirmed that, prior to patching, client devices by Amazon, Apple, Google, Samsung, Raspberry, and Xiaomi, as well as some access points by Asus and Huawei, were vulnerable to Kr00k. These totaled over a billion WiFi-capable devices and access points, at a conservative estimate. Further, many other vendors whose products we did not test also use the affected chipsets in their devices. The vulnerability affects both WPA2-Personal and WPA2-Enterprise protocols, with AES-CCMP encryption."

They said: "Kr00k is related to KRACK (Key Reinstallation Attacks), discovered in 2017 by Mathy Vanhoef, but is also fundamentally different. In the beginning of our research, we found Kr00k to be one of the possible causes behind the reinstallation of an all-zero encryption key observed in tests for KRACK attacks. We responsibly disclosed the vulnerability to chip manufacturers Broadcom and Cypress, who subsequently released updates during an extended disclosure period. We also worked with the Industry Consortium for Advancement of Security on the Internet (ICASI) to ensure that all potentially affected parties - including affected device manufacturers using the vulnerable chips, as well as any other possibly affected chip manufacturers - were also aware of Kr00k."

They said: "According to our information, patches for devices by major manufacturers have been released by now." Except, wait a minute, I'll get there in a second. They said: "To protect yourself as a user, make sure you've applied the latest available updates on all your WiFi-capable devices, including phones, tablets, laptops, IoT devices with WiFi, and WiFi access points and routers."

Leo: [Whimpering]

Steve: "As a device manufacturer" - I know. "As a device manufacturer, please inquire about patches for the Kr00k vulnerability directly from your chip supplier." What I got a kick out of was that, despite the researchers' apparently diligent work to prewarn vendors, it appears that some major suppliers were still caught flat-footed. The day after ESET's startling announcement and disclosure presentation during RSA, Cisco announced that it is "working to patch multiple products that are affected by the recently disclosed Kr00k vulnerability in WiFi chips from Broadcom and Cypress." In the case of Cisco, the Kr00k vulnerability affects at least 14 identified so far Cisco products, I mean, including

low-level enterprise and consumer routers. Cisco said that it is "currently investigating its line of products to identify which ones are vulnerable."

Leo: Which ones? Which ones?

Steve: Geez. And so far it found, like, all of them, all 14 that it has looked at because of course they are using the Broadcom chip. So anyway, a clever discovery. You send a dissociate packet, just to a device as it's transmitting, and suddenly you get 32K of unencrypted data. If you time your disassociation right, that'll be the username and password to a website that you're logging onto over a secure connection, presumably. So, yikes.

Leo: So what are they - they're able to get my traffic.

Steve: They're able, yes, it decrypts your traffic between your device and the access point. It's worth noting, if the traffic is itself encrypted, that is, if you have an HTTPS connection, then it's the WiFi wrapper that is decrypted, not the interior.

Leo: So you're still secure. Or if you're using a VPN. This would be a good argument for getting that VPN fired up.

Steve: Yes, it would. So you're still secure if you're over a VPN. You're also secure, for example, this would be like the decryption of a non-encrypted WiFi. So if you were using open WiFi in a Starbucks...

Leo: Be just like that, yeah.

Steve: Then, yeah, exactly. Then there's no encryption on the WiFi network. Therefore you're relying upon the interior, you know, your own encryption, HTTPS encryption and, like, internal communications encryption. So it's not that big a deal. But for IoT devices, who knows what sort of mischief you can get up to by timing an IoT device's communication with its network, which is assuming WiFi encryption which no longer exists. So again, this gets added to the bag of tricks that the bad guys will have and will be able to deploy over time.

Leo: Wow. Wow, wow, wow.

Steve: Yup. It's good we're finding these things, but we really do seem to be creating them as quickly as we're eliminating them. And older things will never get updated. They will be Kr00ked for life.

Leo: Kr00ked.

Steve: They will be Kr00ked.

Leo: You're Kr00ked. All right. Well, thank you for filling us in on it, anyway. Sounds like there will be patches to most things eventually, or...

Steve: They've already, well, see, so I'm sure that Google has already got their Nexus devices patched. I know Apple has got theirs patched.

Leo: Okay. So, and if you're using a - this is just a good argument for using a VPN.

Steve: It is, yup.

Leo: Yeah. And that's why we're glad HTTPS is so widespread.

Steve: Yes. More so than ever. We just wish our certificates were as secure as they could be. Maybe we'll get that next.

Leo: Staple your certificates, kids.

Steve: In the meantime, take a little drop of sunshine, a little drop of Vitamin D.

Leo: Only one or two polls have been closed. You're going to run, and you're going to go watch the TV and get all the results. I will follow - oh, already I can give you some results.

Steve: No, don't.

Leo: Nothing surprising. No, no, save it. No spoilers. Spoiler alert. New York Times...

Steve: The market, the stock market was suffering again today, as well.

Leo: Lisa said just don't look. Just don't look.

Steve: The New York Times is saying what again?

Leo: You want to be surprised, don't you?

Steve: Okay, okay.

Leo: They've already called a couple of states, let's put it that way.

Steve: I'm all into bonds. I have nothing in the stock market.

Leo: Oh, that's even worse.

Steve: No, because I'm not selling them. I'm letting them - I'm taking them to maturity.

Leo: Yes. I am, too. I'm holding. I'm a buy-and-hold kind of guy.

Steve: Yup, yup.

Leo: Steve Gibson - and, by the way, I buy and hold mutual funds. I don't buy and hold individual stocks because, you know, obviously for the work I do I don't want to be exposed.

Steve: Don't want your finger on the scale.

Leo: No. Steve Gibson is at GRC.com, where he does keep his finger on the scale, but in a good way. He keeps his coffee on the scale, too. He's got a Pringle's can full of coffee. He's putting a Starbucks wrapper around it just to fool you. You could use that as a WiFi antenna. GRC.com is his website. That's where you'll find SpinRite, the world's finest hard drive recovery and maintenance utility. And by the way, if you buy it now, you'll get the upgrade to SpinRite 6.1 free. It'll be part of the deal.

Steve: Before. Before everybody else.

Leo: Yeah, you get to test it, yup. GRC.com. While you're there, of course there's the 16Kb version of the show for bandwidth-impaired folks. There's the 64Kb version for people who like full sound - full, rich sound. There's also transcripts for people who like to read along. That's a nice thing to have. GRC.com. He's on the Twitter at @SGgrc. He takes DMs there, so if you have a question or a comment or suggestion, you can leave it there.

Steve: So much good stuff there. Oh, my god, I need to spend more time there. And every time I look, I'm like, ooh, that's good to know.

Leo: Oh, I'm glad to hear that.

Steve: So thank you, all listeners.

Leo: People are giving you some good stuff. We have audio and video of the show. If you're crazy enough to want to watch, you can do that at TWiT.tv/sn. Or subscribe. Best thing to do would be subscribe because it's an RSS feed, so that way you'll just get it automatically the minute it's available.

We do this show Tuesdays about 1:30 Pacific, 4:30 Eastern, 21:30 UTC. Next week it will be at a different time because...

Steve: What?

Leo: Well, we're going - no, you and I will be the same. But we are springing forward on Sunday.

Steve: Finally.

Leo: Finally.

Steve: Yay, we'll have summertime, yeah.

Leo: Summertime, summertime. So we are springing forward. But that means - but UTC never moves. So instead of 21:30 we'll be at 20:30 UTC.

Steve: We should stop that. We should just leave it in summertime mode.

Leo: So freaking confusing. So, but I'm just glad that the farmers are saving kerosene. That's the matter right here. That's all we care about. Save your kerosene for the war effort. Yes, watch us live at TWiT.tv/live. You can then chat along with the conversation on irc.twit.tv. I think that's all I need to say. You go watch TV. I will be gone tomorrow because we're going to St. Louis. Meet me in St. Louis for WWT, TWiT.tv/blog for more information.

Steve: Stay healthy.

Leo: I will stay - I'm going to chugalug some Vitamin C before I go, and Vitamin D, too, before I go.

Steve: And they're saying now you just bump elbows. You should not [crosstalk] handshake.

Leo: No touching. No touching. No bad touch. And don't wear a face mask because the doctors need those. And wash your hands a lot. I've got it all. Thank you, Steve.

Steve: Good.

Leo: We'll see you next week on Security Now!.

Steve: Bye.