Security Now! #756 - 03-03-20

Kr00k

This week on Security Now!

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**."

Huh???

"SandSara" on Kickstarter: https://grc.sc/sand

Security News

A significant milestone for Let's Encrypt: The free service's **BILLIONTH** certificate issued.

As we know, Netscape initially defined SSL and created HTTPS in 1994. But its adoption is a repeat of the classic Internet tale of adoption reluctance:

Month	% HTTPS
Aug 2018	51.78%
Feb 2018	38.42%
Aug 2017	30.78%
Feb 2017	19.96%
Aug 2016	13.76%
Feb 2016	9.39%
Aug 2015	6.71%

Even 21 years later, in August of 2015, a mere 6.71% of the Alexa Top Million sites were redirecting their users to HTTPS. And, believe it or not, it wasn't until August 2018 -- just a year and a half ago -- that the number finally broke across the 50% mark, coming in at 51.78%.

But there's no doubt that a combination of factors merged to make this happen. For one thing, the creation of "Let's Encrypt" eliminated the financial and "annual hassle" arguments against switching a website to encryption. We began talking about Let's Encrypt when it was first publicly announced back in November of 2014. A non-profit effort from the Internet Security Research Group (ISRG), its development work was sponsored by the Electronic Frontier Foundation (EFF), Cisco, Facebook, Google, the Internet Society (where the Internet Engineering Task Force (IETF) lives), Mozilla, and the French cloud service provider OVH.

This work brings us ACME -- the Automated Certificate Management Environment -- which, as we know, automates the proof of domain ownership and the issuance and installation of certificates, each having a maximum 90-day lifetime. But even with certificates being free, self-installing and self-renewing, nearly three years after its introduction, only 58% of webpages were being delivered over HTTPS. Today we are at 81% But that's still only 4 out of every 5.

The other impetus which is just now beginning to take hold are our browsers' incrementally strengthening shaming of any and all non-HTTPS webpages. This will only increase as we move forward and as excuses for not encrypting continue to lose credibility. And when you consider that Google gives HTTPS sites a higher Search Ranking than the otherwise equivalent HTTP site, the incentive for those remaining 19% of sites to switch to HTTPS becomes very real.

So... a billion certs, and counting. That's a true milestone for the Let's Encrypt effort. Though, as we've been noting, the automated issuance of a certificate for **any** domain, including "PlayPal.com", changes the assertion being made by that certificate. Since no one but we geek techies 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. It's as likely as not to be a forged spoofed site. But, hey, at least it'll be a secure forgery!

I know that I may seem to be harping on this issue. So allow me to provide a bit of perspective before we move on to our next topic. Research conducted back on on March 20th, 2017 revealed that Let's Encrypt had issued 15,270 "PayPal" certificates -- certificates either containing the term "PayPal" or some visual look alike phrase. The actual "PayPal" certificate is an extended validation EV cert obtained, as is mine, from DigiCert. Imagine how PayPal feels about having 15,270 lookalike certificates issued to secure spoofing domains created to confuse their users. And this is why I believe it would be useful and meaningful to users to have our web browsers indicate when a website is being protected by a certificate what was obtained by an organization that some human took a few minutes to verify and vet.

Facebook, Google, and Twitter threaten to leave Pakistan altogether - and together
The government of Pakistan has published some proposed stringent censorship rules governing online content that have Facebook, Google, and Twitter threatening to pull up stakes and go entirely dark throughout the country of Pakistan. Known as the Pakistan "Citizens Protection (Against Online Harm) Rules, 2020.":

https://mcusercontent.com/3db897db1506081dc74dd704d/files/39198ebc-1e0f-46b8-98f5-b88eb7ab8bee/CP_Against_Online_Harm_Rules_2020.pdf

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."*

Furthermore, 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." The rules also give the government the right to block a social network if they refuse to comply or impose fines of up to five hundred million rupees -- which, in this case is more than $50. It's approximately $6.9 million USD.

Those fighting for free speech online argue that such wide-reaching powers are designed to curb free speech and impose censorship. And Facebook, Twitter, and Google appear to agree. Those three organizations are part of the Asia Internet Coalition (AIC), a trade association discussing issues surrounding Internet innovation and regulation in the region. In response to Pakistan's proposed rules, the AIC 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, 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 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 risking the "crippling" of Pakistan's emerging digital economy. Whoops!

The New York Times, which first reported on this, noted that by threatening to leave altogether, the companies may be attempting to apply pressure 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.

However, Pakistan is not alone. Last year, India proposed a set of rules in the same vein, prompting the same concerns over free speech. India is expected to publish these guidelines soon.

So, we have a bit of a mess on our global hands. Our Western social media companies were created in the comparatively permissive and open United States where we enjoy a great deal of freedom. But even here in the US we're seeing increasing trouble with online hate speech, counter-factual content posted and posing as factual, and the fact that "news" is increasingly available from unvetted sources. And then, of course, there's the encryption problem.

Remember our "friends" at Clearview AI?
This was the company who was scraping facial content from the web and reselling it to unnamed but presumably highly interested 3rd-parties without the knowledge or permission of those whose faces had been captured by their database?

Well... In an interesting turn of events, those 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 (whoopsie!), along with information including the number of searches those customers have made and how many accounts customers had set up.

Recall that the app identifies people by comparing photos to a database of images scraped from social media and other sites. It first rose to our attention when a New York Times investigation into the software company earlier this year revealed its activities. At that time The Times revealed that Clearview had quietly sold access to faceprints and facial recognition software to more than 600 law enforcement agencies across the US, claiming that it can identify a person based on a single photo, reveal their real name and far more.

Within a few weeks of the Times article, Clearview was being sued in a potential class action lawsuit that claims the company amassed the photos out of "pure greed" to sell to law enforcement, thereby violating the nation's strictest biometrics privacy law – Illinois' Biometric Information Privacy Act (BIPA) – where **any** biometric data, including a person's own face, cannot be captured and used without their explicit permission.

Learning of this, Senator Edward Markey called Clearview AI a “chilling” privacy risk. And since then, Facebook, Google, YouTube, Microsoft and Twitter have all sent cease-and-desist letters to Clearview AI.

So who exactly has been dipping into the Clearview well? Not only the expected law enforcement agencies like ICE (Immigration and Customs Enforcement) and the US DOJ (Department of Justice), people at the FBI, Customs and Border Protection and Interpol...

But also, according to reporting by BuzzFeed News, AT&T, Verizon, T-Mobile, Best Buy, Eventbrite, Las Vegas Sands, Coinbase, Bank of America, Walmart, Kohl's and Macy’s.

The privacy geeks are naturally all freaked out. Nathan Freed Wessler, a staff attorney with the American Civil Liberties Union (ACLU) said that “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 and 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 emailed 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. (Notice that, as usual, everyone is hiding behind children.)

AT&T said it's not a client of Clearview AI. Best Buy denied ever using or planning to use Clearview AI. Bank of America 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.

Coinbase said it hasn't yet made a commitment to use Clearview AI. In an emailed statement a Coinbase spokesperson wrote: “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. We have 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 (in 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. Yikes!

Evan Greer, deputy director of Fight for the Future, said in a statement: "This is a crisis for our democracy. Lawmakers need to get off their butts, do their jobs, and pass legislation to ban the use of facial recognition surveillance not just by government agencies but by corporations too."

Business appears to be brisk for Clearview AI and facial recognition through social media content scraping and analysis. What a brave new world.

Swiss government submits criminal complaint over CIA Crypto spying scandal
So, this whole Crypto AG thing has predictably blown up with apparently well-founded allegations that US and German intelligence deliberately implemented backdoors in Crypto AG's systems to eavesdrop on governments worldwide.

The Swiss complaint focuses upon Operation Rubicon, as a result of a recent investigation by the Washington Post, ZDF (a German public-service television broadcaster), and SRF News (a Swiss Radio Station) into the Swiss company Crypto AG.

Crypto AG is a seller of encoded and encrypted devices deemed suitable -- and secure enough -- for confidential government communications. It is estimated that over 100 governments worldwide have been counted as Crypto AG clients over the course of decades. Rumors concerning the CIA and its German counterpart BND being able to crack these devices have been around for some time, and now the recent inquiry -- which reveals that Crypto AG was, until recently, actually owned by these authorities -- claims that the agencies deliberately introduced backdoors and weaknesses in products sold by Crypto AG to intercept and eavesdrop on users.

The report says that governments were paying "good money to the US 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," which further suggests that communication records may have been shared between Five Eyes members -- being the US, UK, Canada, Australia, and New Zealand.

Being a Swiss company, Crypto AG covert operations has shaken Switzerland due to its long-term standing as a neutral country in political matters, a reputation that it has defended for years.

In 2018, Crypto AG split into two companies: 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. CyOne AG maintains it is independent of its international counterpart.

The Swiss government has appointed Niklaus Oberholzer, a former supreme court judge, to investigate the matter. Oberholzer's 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 (SECO).

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 in the country knew about the surveillance practices in the hopes of mitigating damage caused to Switzerland's neutral position. A local Swiss publication reports that the complaint refers to the Swiss Goods Control Act, legislation designed to control product sales in both civilian and military capacities. Those who fall foul of the act, such as by providing false or incomplete information on petition for a sales license, could face a decade behind bars and fines of up to five million Swiss francs. The complaint alleges that the State Secretariat for Economic Affairs was hoodwinked and misled into permitting the sale of Crypto AG goods and would not have granted a license if the authority knew of the surveillance scheme. (Yeah, no kidding.)
"Ghostcat" -- (Apache) Tomcat Users: Update NOW!

We have another serious problem in a widely deployed web server.

Tomcat is an open source implementation of the Java Servlet, JavaServer Pages, Java Expression Language and Java WebSocket technologies. So, Tomcat provides a "pure Java" HTTP web server environment for hosting Java-based web applications. It started off 22 years ago, back in 1998 as a reference implementation of a Java servlet created by James Duncan Davidson who was a software architect at Sun Microsystems. Since then it has been evolving steadily across a series of minor and major releases.

And, targets will not be difficult to locate. Shodan lists more than 890,000 Tomcat servers currently reachable over the Internet, and the similar BinaryEdge service has located more than 1 million.

Ghostcat is a high-risk file read/include vulnerability tracked as CVE-2020-1938 and present in the Apache JServ Protocol (AJP) of Apache Tomcat between versions 6.x and 9.x. That's ALL VERSIONS of the Apache Tomcat server released during the past 13 years.

Critical flaw that can lead to server takeover

The Tomcat developers, in a vain attempt to put the best possible face on this disaster write: "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."

Researchers at the Chinese security firm Chaitin Tech, who discovered the bug explained that after successfully exploiting an unpatched Tomcat server "an attacker can read the contents of configuration files and source code files of all webapps deployed on Tomcat. In addition, if the website application allows users to upload files, an attacker can first upload a file containing malicious JSP script code to the server (the uploaded file itself can be any type of file, such as pictures, plain text files etc.), and then include the uploaded file by exploiting the Ghostcat vulnerability, which will result in remote code execution."

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 JBoss Web Server (JWS) and JBoss Enterprise Application Platform (EAP).

And, not at all 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...

https://cwiki.apache.org/confluence/display/TOMCAT/PoweredBy

... the cyber threat intelligence firm "Bad Packets" tweeted on Saturday: “mass scanning activity targeting this vulnerability has already begun. PATCH NOW!”

Revisiting OCSP Must Staple

I was thinking about OCSP Must Staple. I hadn't looked at it for a long time. But I was wondering why it might NOT be the right solution. And 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 could still be used because the malicious server would not ask browsers to require stapling. So I double-checked.

I found security researcher Scott Helme's clearly-written page on OCSP Must-Staple: https://scotthelme.co.uk/ocsp-must-staple/ His page starts out saying:

> Revocation checking is broken and has been for some time. Whilst some vendors have sort of worked around this with proprietary solutions, there is little that the smaller sites can do. OCSP must-staple to the rescue!"
>
> 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 had 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 to the CA. 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. Thus, we finally arrived at OCSP Must-Staple.
>
> ***OCSP Must-Staple***
> Setting up OCSP Must-Staple is fairly easy as it's simply a flag that needs to be set by your CA in the certificate they generate for you. This flag instructs the browser that the certificate must be served with a valid OCSP response or the browser should hard fail on the connection.

So, in other words, the good news is, it's not up to the server. The "must staple" flag is bound into the CA-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 with a one-year expiring certificate.

The only reason I can see for NOT taking the "must staple" path would be that we're not yet ready for its deployment. So I did some digging and found a year-old assessment. https://blog.apnic.net/2019/01/15/is-the-web-ready-for-ocsp-must-staple/

	Desktop Browsers											Mobile Browsers				
	Chrome 66			Firefox 60			Opera		Safari	IE	Edge	Safari	Chrome		Firefox	
	OS X	Lin.	Win.	OS X	Lin.	Win.	OS X	Win.	11	11	42	iOS	iOS	And.	iOS	And.
Request OCSP response	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Respect OCSP Must-Staple	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓
Send own OCSP request	✗	✗	✗	-	-	-	✗	✗	✗	✗	✗	✗	✗	✗	✗	-

As of January 2019, only Firefox 60 of all browsers respected OCSP Must Staple. Chrome 66 did not. Opera did not. Safari did not. IE11 did not. Edge 42 did not.

And the last update on the server side indicated that Microsoft's IIS supports OCSP-Stapling properly but that 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. But because no one has pushed the open source servers nor the end-user's 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. You really would think that Google, with their proven ability to guide the industry due to the strength of Chrome and Chromium, would be leading the way here. After all, their browser's revocation with the CRLSET solution remains, as we know, the worst of all. But that hasn't happened yet.

It really looks as though this is where we as an industry should focus. Once our web servers are up to speed and all of our browsers honor must-staple, then nothing would prevent CAs from issuing 10-year certs that are FAR more secure than anything we have today. And browser behavior could be adaptive. Apple could setup Safari to limit the total lifetime of all non-must -staple certificates to one year, but to allow any bearing the "must staple" flag to have any lifetime that it wishes. Nothing is preventing this other than our own laziness.

Miscellany

All this recent news about the COVID-19 Coronavirus pandemic had me thinking about Security Now Episode #209, which we recorded mid-August of 2009. As that episode describes itself: "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: Vitamin D."

So, that was 10 and a half years ago. And I recall that afternoon clearly, Leo."...

Google: "GRC Vitamin D":
https://www.grc.com/health/vitamin-d.htm

SpinRite

I've started working on the AHCI driver for SpinRite. I already have low-level BIOS-bypassing hardware drivers in place for all pre-AHCI technologies, traditional IDE and Legacy SATA controllers. So the final piece of technology, as I mentioned last week, will be 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. 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.

Kr00k

The researchers at ESET named their latest discovery "Kr**00**k" to highlight the **zeroes** in the name... because "zeroes" is what this latest discovery of a very serious WiFi vulnerability affecting more than a billion devices is all about.

The vulnerability, which I'll explain in a moment, affects anything not yet patched containing the most popular WiFi chips in the industry, which are manufactured by Broadcom and Cypress, and include 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. All are, or were, vulnerable to Kr00k. In other words, all consumer devices, access points, routers, and IoT gadgets -- anything using WiFi with chips from Broadcom or Cypress. So, we know that weighs in at more than a billion WiFi-capable things. And since the Cypress chips are popular among IoT devices, their numbers are untold.

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. The bad news is, 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.

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 fear of the unknown. It's based upon solid practical evidence. Given everything we have learned about the desires and motivations of hackers and state actors alike, I guarantee you that somewhere -- or more likely at many somewheres -- people are methodically compiling these increasing number of things that, for some large swath of the device population, are never going to be fixed. For some time now I've been characterizing security as "porous." But it's becoming more so than I expected.

Like all of the worst and most powerful attacks, this one is not that complicated. You only have to know about it.

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 the shared cryptographic key given a shared secret. For example, our home router knows our WiFi network's password. So we give the same password to our laptop computer and because each now knows the same shared secret they are able to negotiate a session key which will be used to encrypt the communications.

We've talked about attacks on WiFi in the past. Being radio, denial of service is easy. One brute force approach would be to jam the receiving end with RF noise. But a much more subtle attack is made possible by the fact that the 802.11 WiFi connection management packets are never encrypted. So it's possible to deny any unwitting WiFi user the use of the nearby WiFi simply by spoofing and sending "disassociate" packets to their laptop. The "disassociate" command will be immediately obeyed, and they will drop off the net. However, their endpoint will then immediately start working to reassociate itself with the access point, so this process would need to be repeated.

But the clever Kr00k vulnerability occurs as a side-effect of that disassociation: Inside the WiFi firmware is the negotiated encryption key, and a sizeable -- like 32k -- transmission buffer. As one would hope and want, when a WiFi session is disassociated, the transient session encryption key they were using is immediately and proactively zeroed. It is wiped and set to all zeroes. But, unfortunately, that 32k buffer is not also 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 WPA-2 AES CCMP cipher mode operates, 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.

https://www.welivesecurity.com/wp-content/uploads/2020/02/ESET_Kr00k.pdf

ESETs disclosure White Paper says:

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 Wi-Fi chips by Broadcom and Cypress that haven't yet been patched. These are the most common Wi-Fi chips used in contemporary Wi-Fi-capable devices such as smartphones, tablets, laptops, and IoT gadgets. Not only client devices but also Wi-Fi access points and routers with Broadcom chips were affected by the vulnerability, thus making many environments with unaffected or already patched client devices vulnerable anyway.

[In other words, yeah... you have an Apple iPhone that Apple has already patched, but what about your old router? And what about the router at Starbucks or the airport or on the plane?]

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 Wi-Fi-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.

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 aware of Kr00k.

According to our information, patches for devices by major manufacturers have been released by now. To protect yourself, as a user, make sure you have applied the latest available updates on all your Wi-Fi-capable devices, including phones, tablets, laptops, IoT devices with Wi-Fi, and Wi-Fi access points and routers. As a device manufacturer, please inquire about patches for the Kr00k vulnerability directly with your chip manufacturer.

Despite the researcher's apparently diligent work to pre-warn vendors it appears that some major suppliers were caught flat-footed. The day **after** ESET's startling announcement/ presentation at 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 Cisco products. Cisco said that it was "currently investigating its line of products to identify which ones are vulnerable." So far it has identified 14.