Security Now! #900 - 12-06-22

LastPass Again

This week on Security Now!

This week we answer a few questions: What if an Australian company doesn't secure their own network? Has Ireland NOT levied fines against any major Internet property owned by Meta? What's in REvil's complete dump of Australia's Medibank data disclosure? We finally answer the question: Is nothing sacred? (It turns out it’s not rhetorical.) Also, whose root cert just got pulled from all of our browsers, and how did a handful of Android platform certs escape? What US state has banned all use of Tik-Tok? What country is prosecuting its own ex-IT staff after a breach? How has memory-safe language deployment actually fared in the wild? Are last August's BlackHat 2022 videos out yet? And which brand of IoT security camera do you probably NOT want to use or purchase? Which podcast had the most amazing guest last week? What happened when SpinRite was run on an SSD? And what does LastPass's announcement of another hacker intrusion mean for it and its users? Answers to those questions and more coming your way during this week's Security Now! podcast.

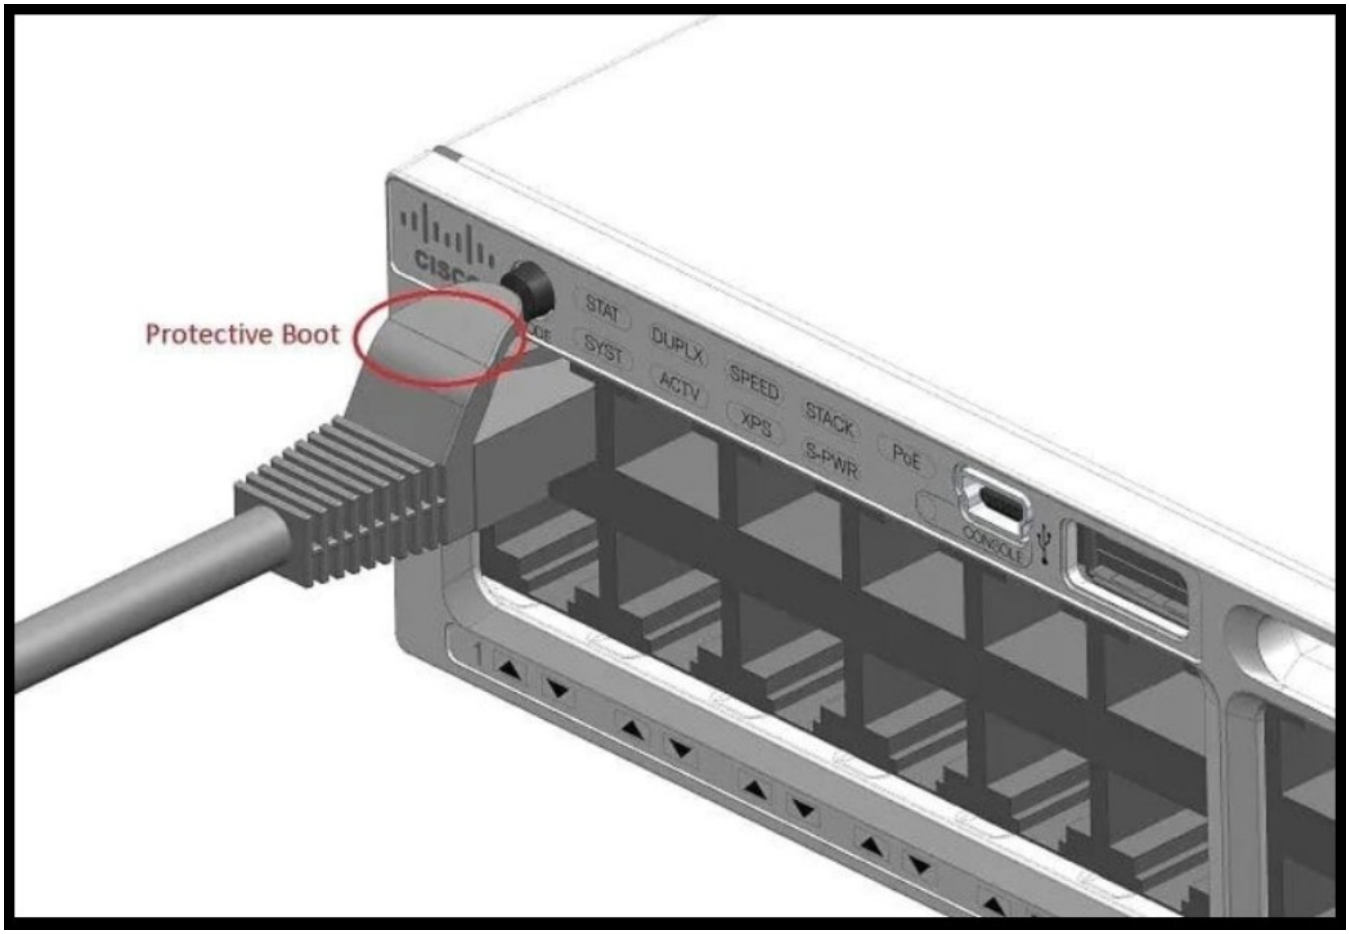

If you've ever messed up a dimension or a hole position on something you're building, don't be too hard on yourself. At least you're not the Cisco design engineer who caused an entire product line recall by placing the mode button directly above an RJ45 port. That button resets the switch to its factory settings when it's held down.

Security News

Don't mess with Australia

Continuing our recent Australia watch: Recall from last week that a recent cybersecurity country ranking published by MIT gave Australia in the #1 spot for Cyber Defense, followed by, in decreasing order, the Netherlands, South Korea, the US and Canada. We also recently covered Australia's proactive declaration of CyberWar which they will be waging against the world's perpetrators of cybercrime – not waiting for them to commit another crime. It would seem that those high profile cyberattacks on Optus, Telstra, Medibank, Woolworths, and EnergyAustralia really woke the sleeping bear and galvanized Australia into action. They've decided to go active.

Last week saw another facet of this campaign, with the creation of new legislation to replace Australia's creaky 34-year old Privacy Act of 1988. The new legislation ups the ante when Australia's own internal attack targets are willfully negligent. The legislation bears the cumbersome name "Privacy Legislation Amendment (Enforcement and Other Measures) Bill 2022." It grants the Office of the Australian Information Commissioner (OAIC) the power to levy hefty fines on companies — and not only Australian companies — that ignore security best practices to needlessly expose their customers' data through cybersecurity breaches.

Under this bill, which is expected to receive royal assent shortly to place it into force, companies that fail to safeguard user data face fines of up to the greater of AUS $50 million or 30% of the company's adjusted revenue. The existing fines impose only a AUS$2.22 million fine as a result of security breaches. So we're going from $2.22 million to $50 million OR 30% of revenue if 30% of profit is more than $50 million.

Steep as this is, the updating of Australia's antiquated legislation has, not surprisingly, been greeted with positive feedback from Australian cybersecurity experts, who view it as the incentive needed to get local companies to pay the attention that's needed to the state of their IT systems. Given the historical reticence that we keep seeing to getting ahead of this problem, I can't see any other way to bring about the changes we need. It does still feel wrong, however, for the suppliers of the too-often buggy systems upon which enterprises rely continue to be held harmless.

I mentioned before that this law wasn't applicable only to Australian companies. According to the bill's text, its provision and fines will also apply to any non-Australian company who is doing business in Australia, even if they are headquartered outside of Australia. And THAT promises to prove interesting!

Facebook / Meta fined by Ireland

While we're on the topic of fines, Ireland's data protection agency fined Meta €265 million due to Facebook's data breach a year and a half ago in April of 2021. The Irish Data Protection Commission said that Meta failed to safeguard its Facebook platform from data scraping, which allowed a threat actor to compile details on more than 530 million Facebook users. This data was later sold on an underground cybercrime forum. Facebook told TechCrunch that following the incident they rolled out protections to detect scraping operations. Since Ireland had previously fined Meta's Instagram €405 million in September 2022 and Meta's WhatsApp €228 million in September 2021, this rounds out the fines so that Ireland's data protection agency has now fined each of Meta's three main platforms.

REvil's full Medibank dump:
Remember that Medibank was one of the several ransomware embarrassments that Australia-based organizations recently suffered. Well, Medibank stood up to the REvil gang, refusing to pay or to buckle under to REvil's extortion, threatening to release the entire contents of what had been illegally stolen from Medibank. Medibank is Australia's largest health insurer and what's known is that the significant personal data for Medibank's 9.7 million clients was stolen during REvil's original intrusion and data exfiltration.

So, last Thursday Medibank released a lengthy statement which began:

> *We are aware that stolen Medibank customer data has been released on the dark web overnight. We are in the process of analyzing the data, but the data released appears to be the data we believed the criminal[s] stole. Unfortunately, we expected the criminal[s] to continue to release files onto the dark web.*
>
> *While our investigation continues, there are currently no signs that financial or banking data has been taken. And the personal data stolen, in itself, is not sufficient to enable identity theft and financial fraud. The raw data we have analyzed today so far is incomplete and hard to understand.*

A bit later in their released statement, quoting Medibank's CEO David Koczkar, they wrote:

> *Anyone who downloads this data from the dark web, which is more complicated than searching for information in a public internet forum, and attempts to profit from it is committing a crime.*
>
> *The Australian Federal Police have said law enforcement will take swift action against anyone attempting to benefit, exploit or commit criminal offenses using stolen Medibank customer data. We continue to work closely with the Australian Federal Police who are focused, as part of Operation Guardian, on preventing the criminal misuse of this data.*
>
> *Again, I unreservedly apologize to our customers. We remain committed to fully and transparently communicating with customers and we will continue to contact customers whose data has been released on the dark web.*

At the end of the statement Medibank enumerates the sobering details of what they believe the REvil gang both obtained and did not obtain. On the daunting did obtain side are:

- *The name, date of birth, address, phone number and email address for around 9.7 million current and former customers and some of their authorized representatives. This figure represents around 5.1 million Medibank customers, around 2.8 million ahm customers and around 1.8 million international customers*
- *Medicare numbers (but not expiration dates) for ahm customers.*
- *Passport numbers (but not expiration dates) and travel visa details for international student customers.*
- *Health claims data for around 160,000 Medibank customers, around 300,000 ahm customers and around 20,000 international customers. This includes service provider name and location, where customers received certain medical services, and codes associated with diagnosis and procedures administered.*

- *Additionally, around 5,200 My Home Hospital (MHH) patients have had some personal and health claims data accessed and around 2,900 next of kin of these patients have had some contact details accessed.*
- *Health provider details, including names, provider numbers and addresses.*

On the flip side they said that REvil:

- *Did not access primary identity documents, such as drivers' licenses, for Medibank and ahm resident customers. Medibank does not collect primary identity documents for resident customers except in exceptional circumstances.*
- *Did not access health claims data for extras services (such as dental, physio, optical and psychology).*
- *Did not access credit card and banking details.*

What most caught my attention was their statement that "*The raw data we have analyzed today so far is incomplete and hard to understand.*" What occurs to me is that a raw unorganized dump of data concerning 9.7 million current & past customers is far less actionable than an organized searchable online database containing the same information. In other words, it's almost entirely the structure of the data that gives it meaning and makes it useful. If REvil grabbed raw data files without formatting templates and indexes into the data, if the database is highly relational and deeply depends upon its relations and indexes for it to be pulled together into something coherent, then the release of a massive blob of raw and disorganized data, where there's nothing to make clear what piece goes with what, might be much less damaging than it at first appears.

Is nothing sacred?

Okay. So I couldn't resist giving this short bit of news the heading "Is nothing sacred" because the official website of The Vatican was pushed offline last Wednesday by a DDoS attack which was carried out by pro-Russian hacktivists. As CNA, the Catholic News Agency pointed out, the attack came a day after Moscow criticized Pope Francis's latest condemnation of Russia's invasion of Ukraine. So, no... nothing is sacred.

Mozilla yanks a (no longer) trusted root

As we know, because we've covered this through the years, web browsers are extremely reticent to remove root certificates from their trusted root stores because doing so immediately renders invalid and not trusted any and all outstanding certificates which have been previously signed by that certificate authority's matching private key. This effectively puts the certificate authority out of business overnight. This has happened a few times since the start of this podcast, and it's always interesting.

In this case, the certificate authority in question is an apparently shady Panamanian firm called TrustCor. Nearly a month ago, long simmering questions about TrustCor were brought to a boil by a piece in the Washington Post whose headline was: *"Mysterious company with government ties plays key internet role. TrustCor Systems vouches for the legitimacy of websites. But its physical address is a UPS Store in Toronto."* That'll get your attention.

Here's just a sampling of the juicy bits from the The Washington Post's reporting:

> *The company's Panamanian registration records show that it has the identical slate of officers, agents and partners as a spyware maker identified this year as an affiliate of Arizona-based Packet Forensics, which public contracting records and company documents show has sold communication interception services to U.S. government agencies for more than a decade.*
>
> *One of those TrustCor partners has the same name as a holding company managed by Raymond Saulino, who was quoted in a 2010 Wired article as a spokesman for Packet Forensics. Saulino also surfaced in 2021 as a contact for another company, Global Resource Systems, that caused speculation in the tech world when it briefly activated and ran more than 100 million previously dormant IP addresses assigned decades earlier to the Pentagon. The Pentagon reclaimed the digital territory months later, and it remains unclear what the brief transfer was about, but researchers said the activation of those IP addresses could have given the military access to a huge amount of internet traffic without revealing that the government was receiving it.*

And recall that we talked about this weird event at the time, noting how odd it was that this previously dormant and DoD-reserved block of IPv4 addresses was suddenly being routed, and tied to some random company who no one had ever heard of. The Post continues...

> *TrustCor's products include an email service that claims to be end-to-end encrypted, though experts consulted by The Washington Post said they found evidence to undermine that claim. Researchers said that a test version of the email service also included spyware developed by a Panamanian company related to Packet Forensics. Google later banned all software containing that spyware code from its app store.*
>
> *A person familiar with Packet Forensics' work confirmed that it had used TrustCor's certificate process and its email service, MsgSafe, to intercept communications and help the U.S. government catch suspected terrorists.*
>
> *Speaking on the condition of anonymity to discuss confidential practices, the person said, "Yes, Packet Forensics does that."*

And come on. The name "Packet Forensics" should be an obvious enough 'tell' about the company's intentions. Remember: Any device that's holding a certificate which is able to sign other end certificates is thereby able to intercept any and all TLS-secured traffic bound for any remote web server. It accepts the connection, examines the domain being requested, creates and signs a TLS certificate on the fly and returns it to the browser.

In this case, so long as all web browsers contained the TrustCor CA root certificate, they would happily accept the on-the-fly signed certificate. So the connection to the intercept middlebox would be encrypted where the middlebox would decrypt the TLS data for completely in-the-clear analysis. The middlebox would then initiate its own connection to the actual destination server so that its interception was invisible... while it continued to surveil all of the intercepted browser's traffic.

Mozilla currently recognizes 166 root certificate authorities, but no longer the three from TrustCor. Our really long time listeners may recall that episode which followed my chance discovery of what was then the explosion in certificate authorities in Firefox's root store. The last time I had looked before, there were like 7 or 8 certificate authorities total. Now there appeared to be hundreds. And I said at the time, this is inherently not good. **All** web browsers are trusting **any** certificates signed by the owners of **any** of these root certs. It makes for an unstable system.

In fairness, things have gone much better than I expected. The industry has been amazingly effective at policing itself and the events of trusted root store abuse have been very few and far between. It is an obvious privilege to be granted Certificate Authority rights. It's a license to print money, but only so long as the certificate authority's signature means something.

The Washington Post's story is not behind any paywall, and it reads like someone's imaginative fiction, while being well researched and backed by facts. I have a link in the show notes for anyone who's curious to know more, and also a link to the Google groups' discussion among all of the industry participants, including TrustCor's rebuttals, which resulted in this outcome:

https://www.washingtonpost.com/technology/2022/11/08/trustcor-internet-addresses-government-connections/

https://groups.google.com/a/mozilla.org/g/dev-security-policy/c/oxX69KFvsm4

Android Platform Certs Escape

While we're on the subject of crucial certificates, and certificate management, last Wednesday the 30th, an internal Google report, originally filed on November 11th, was made public. And on Friday, the security firm Rapid 7 pulled the pieces together for their report.

Google's report is titled: *"Platform certificates used to sign malware."* And under Technical Details they explain: *"A platform certificate is the application signing certificate used to sign the "android" application on the system image. The "android" application runs with a highly privileged user id - android.uid.system - and holds system permissions, including permissions to access user data. Any other application signed with the same certificate can declare that it wants to run with the same user id, giving it the same level of access to the Android operating system."* In other words, it's a full penetration of Android security.

Digging deeper, we find that the Android Security Team discovered several malware samples in the wild that were signed by platform certificates used by major vendors like Samsung, LG, MediaTek, and Revoview. After discovering the incident, the Android Security Team worked with the affected companies to revoke and rotate the leaked platform certificates.

I liked what Rapid 7 had to say about this, because what they said made a lot of sense about what didn't make a lot of sense about this. Rapid 7 wrote:

On November 30, 2022, a Google report initially filed on November 11, 2022 was made public. The report contained 10 different platform certificates and malware sample SHA256 sums, where the malware sample had been signed by a platform certificate — the application signing certificate used to sign the "Android" application on the system image. Applications signed with platform certificates can therefore run with the same level of privileges as the "Android" application, yielding system privileges on the operating system without user input. Google has recommended that affected parties should rotate their platform certificate. However, platform certificates are considered very sensitive, and the source of these certificates is unknown at this time.

This use of platform certificates to sign malware indicates that a sophisticated adversary has gained privileged access to very sensitive code signing certificates. Any application signed by these certificates could gain complete control over the victim device. Rapid7 does not have any information that would indicate a particular threat actor group as being responsible, but historically, these types of techniques have been preferred by state-sponsored actors. That said, a triage-level analysis of the malicious applications reported shows that the signed applications are adware — a malware type generally considered less sophisticated. This finding suggests that these platform certificates may have been widely available, as state-sponsored actors tend to be more subtle in their approach to highly privileged malware.

Okay. So some low-end adware malware was somehow signed by the closely guarded private keys belonging to some of Android's largest and most reputable vendors. Either those closely guarded private keys escaped, or somehow those still-resident keys were used to sign the malware. Either way, the fact that malware was signed means that something went wrong.

What's weird is that any agency that somehow obtains the ability to get any malware signed by major platform keys is not going to waste that awesome privilege on easily-discovered adware. They would treasure that capability and hold it close, choosing to reserve its use for only highly targeted infiltration specifically so that it is never discovered. Now, thanks to the casual misuse of a collection of certificates, whoever or whatever gained the ability to sign with those certs has almost certainly lost those rights. None of the signatures of those certificates will be trusted. Given what we know, it makes no sense.

South Dakota says: No more Tik-Tok

Okay, here's one for ya... Last week, South Dakota's Governor Kristi Noem signed Executive Order 2022-10, which bans all use of the Chinese social media platform TikTok by state government agencies, employees, and contractors. The Executive Order's news release stated that the order is in response to the growing national security threat posed by TikTok due to its data gathering operations on behalf of the Chinese Communist Party (CCP).

Oh, yeah… ya gotta keep your eye on those commies.

The press release quoted Governor Noem saying: "South Dakota will have no part in the intelligence gathering operations of nations who hate us. The Chinese Communist Party uses information that it gathers on TikTok to manipulate the American people, and they gather data

off the devices that access the platform. Because of our serious duty to protect the private data of South Dakota citizens, we must take this action immediately. I hope other states will follow South Dakota's lead, and Congress should take broader action, as well."

The order took effect immediately and applies to all employees and agencies of the State of South Dakota – no more TikTok for you – including persons and entities who contract with the state, commissions, and authorities or agents thereof.

Wow. I really do wish that I would still be alive in another 100 years to see what the Internet has become by then. Will it have succeeded in pulling the world together? Or will the world's fearful leaders have established borders and regional control just as they have everywhere else?

Albania blames its IT staff

Remember the drama we covered back in July, where Iran retaliated against Albania by melting down their government networks. Then Albania retaliated back by severing diplomatic ties with Iran and sweeping into the just-closed Iranian embassy looking for anything that Iran might not have sufficiently destroyed? Also recall that it turned out that Iran had been rummaging around in Albania's networks since April of 2012, so more than 15 months, undetected.

Well, who was to blame for all of this? It must be someone's fault, right? And we can't blame the vendors of the buggy systems. After all, they provided patches. For some of the problems. Usually. Eventually. Again, we've got to blame someone. So Albania has decided that it was all the fault of the IT staff – so now they're in trouble. Albanian prosecutors have charged and asked for the house arrest of 5 government employees. The prosecutors say the 5 accused failed to apply security updates to government systems and also failed to detect that hackers had been wandering around inside their network as far back as April 2021.

Okay. So maybe the IT guys were seriously negligent. But we know that's not necessarily the case. If I may segue for a moment, a perfect example of Albanian-scale negligence not being necessary is the news that the US Department of Homeland Security's Cyber Safety Review Board recently said that it intends to review attacks carried out by the Lapsus$ extortion group and will publish a report detailing how Lapsus$ managed to bypass a broad range of security measures without the use of advanced malware and managed to breach a large number of high-profile targets including Cisco, Microsoft, Nvidia, Samsung, Uber, Rockstar Games, and others. Those companies are not firing their IT department staff because they recognize that it's possible to do nothing wrong and still be breached.

In Albania's case, it could just as easily have been His Excellency the President of the Republic of Albania who clicked a link in a phishing eMail to invite those crafty Iranian cyberwarriors to come for a visit. And who knows what managerial opposition or budgetary constraints the intrepid five might have faced? IT departments are notoriously under staffed, overworked and unappreciated. And IT people are just like everyone else. There are good ones and there are bad ones. Which are they? Who knows. What I wonder, though, is who they're going to get to fill those vacated jobs now? With the risk of prosecution for attempting to do a job that might be impossible, and knowing what happened to the last guys, I would not be surprised to learn that those IT staff positions were difficult to fill.

Good news on the memory safe languages front:
We do have some encouraging news on the memory-safe languages front thanks to some monitoring the feedback about the production and in-field use of memory-safe languages in Android. Since the August release of Android 13, which was the first Android where a majority of the new code added to the project was written in memory-safe languages including Rust, Java, and Kotlin, Google noted that since shifting its focus to using memory-safe languages, the number of memory safety vulnerabilities reported in the Android OS has dropped to less than half of comparable counts.

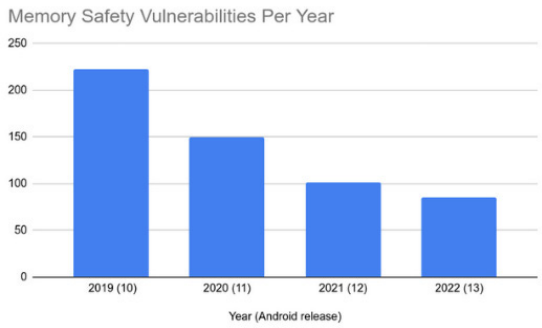

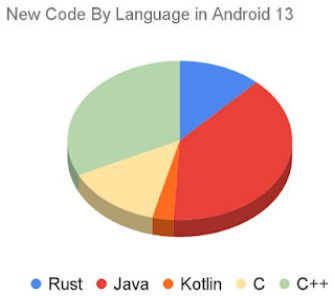

Black Hat USA 2022:
The YouTube videos of the talks from Black Hat USA 2022 which took place, as always, last August, are finally out and available on YouTube:
https://www.youtube.com/playlist?list=PLH15HpR5qRsVKcKwvIl-AzGfRqKyx--zq

Another Chrome 0-day bites the dust
Those using Chrome may have noticed their browser's announcement of an update over the weekend. That squashed another 0-day that had been discovered being used in the wild. It was a type confusion bug in Chrome's JavaScript V8 engine. The vulnerability was discovered by one of the Google TAG researchers. This brings Chrome's year 0-day total to 9.

Anker's Eufy Camera debacle
The Verge's coverage of Anker's Eufy IoT security cameras didn't pull any punches. Their headline read *"Anker's Eufy lied to us about the security of its security cameras."* and their subhead added *"Despite claims of only using local storage with its security cameras, Eufy has been caught uploading identifiable footage to the cloud. And it's even possible to view the camera streams using VLC."* Since I can't improve on The Verge's coverage and reporting. Here's what they wrote:

> *Anker has built a remarkable reputation for quality over the past decade, building its phone charger business into an empire spanning all sorts of portable electronics — including the Eufy home security cameras we've recommended over the years. Eufy's commitment to privacy is remarkable: it promises your data will be stored locally, that it "never leaves the safety of your home," that its footage only gets transmitted with "end-to-end" military-grade encryption, and that it will only send that footage "straight to your phone."*

So you can imagine our surprise to learn you can stream video from a Eufy camera, from the other side of the country, with no encryption at all.

Now, The Verge's coverage of this might seem somewhat harsh. But they then show a snap of the marketing for the camera which makes all of these claims quite clear, which then makes the reality of what Anker is doing somewhat stunning:

This really is shocking. Get ready for lawsuits. The Verge continued:

Worse, it's not yet clear how widespread this might be — because instead of addressing it head-on, the company falsely claimed to The Verge that it wasn't even possible.

On Thanksgiving Day, infosec consultant Paul Moore and a hacker who goes by Wasabi both alleged that Anker's Eufy cameras can stream encryption-free through the cloud — just by connecting to a unique address at Eufy's cloud servers with the free VLC Media Player.

When we asked Anker point-blank to confirm or deny that, the company categorically denied it. "I can confirm that it is not possible to start a stream and watch live footage using a third-party player such as VLC," Brett White, a senior PR manager at Anker, told me via email.

> [The "Senior PR manager at Anker" ... Yep, that's exactly whose opinion you want regarding anything potentially damaging to his employer's reputation.]

But The Verge can now confirm that's not true. This week, we repeatedly watched live footage from two of our own Eufy cameras using that very same VLC media player, from across the United States — proving that Anker has a way to bypass encryption and access these supposedly secure cameras through the cloud.

There is some good news: there's no proof yet that this has been exploited in the wild, and the way we initially obtained the address required logging in with a username and password before

Eufy's website will cough up the encryption-free stream. (We're not sharing the exact technique here.) Also, it seems like it only works on cameras that are awake. We had to wait until our camera's owner pressed a button before the VLC stream came to life.

But it also gets worse: Eufy's best practices appear to be so shoddy that bad actors might be able to figure out the address of a camera's feed — because that address largely consists of your camera's serial number encoded in Base64, something you can easily reverse with a simple online calculator.

The address also includes a Unix timestamp you can easily create, plus a token that Eufy's servers don't actually seem to be validating (we changed our token to "arbitrarypotato" and it still worked), and a four-digit random hex whose 65,536 combinations could easily be brute forced.

"This is definitely not how it should be designed," Mandiant vulnerability engineer Jacob Thompson tells The Verge. For one thing, serial numbers don't change, so a bad actor could give or sell or donate a camera to Goodwill and quietly keep watching the feeds. But also, he points out that companies don't tend to keep their serial numbers secret. Some stick them right on the box they sell at Best Buy — yes, including Eufy.

On the plus side, Eufy's serial numbers are long at 16 characters and aren't just an increasing number. "You're not going to be able to just guess at IDs and begin hitting them," says Mandiant Red Team consultant Dillon Franke, calling it a possible "saving grace" of this disclosure. "It doesn't sound quite as bad as if it's UserID 1000, then you try 1001, 1002, 1003."

I'm reminded of the fact that I don't have a single connected video camera anywhere within my environment, and that Leo's wife Lisa early on intuited the inherent dangers of having unknowable video capture technology – which is what all of this is – lurking around the house.

In a TNO – trust no one – world, the simple though impractical truth is: *"Unless you designed it yourself, you don't know what it does."* And I should add that due to the crazy complexity of the things we design today, even if you did design it yourself, it may not do what you think it does!

Miscellany

An amazing-looking WiFi-6 router... $119
https://www.seeedstudio.com/LinkStar-H68K-1432-p-5501.html

This is moderately random, but not too far afield for this podcast. Everyone knows of my passion for coding. But I predate electronic computers and before computers was electronics. Although coding has taken over, electronics will always be my first love. So in addition to coding, I occasionally do a bit of tinkering, hacking and designing with electronics. At some point in the past some Googling must have taken me to a place called "Seeed Studio" (S e e e d spelled with three 'e's). I purchased something from them and was promptly added to their periodic mailing list.

In this case I don't mind the spam because the mail contains photos of the stuff they're promoting and my jaw spends most of its time hanging down with my mouth open over the insanely low cost of the technology that's currently available from China. It's truly astonishing.

For example, a recent mailing showed the "Seeed Studio XIAO ESP32C3." It's a tiny module about the size of a quarter with 14 electrical connections, 7 per side, and what appears to be two tiny buttons and an LED. It also has a tiny type-C USB connector, presumably for programming. And all of the software for doing so is open course. Its description says: "Seeed Studio XIAO ESP32C3 adopts new RISC-V architecture, supporting both Wi-Fi and BLE wireless connectivities. For Internet of Things applications, you will find it is flexible and suitable for all kinds of IoT scenarios." I was curious so I looked into the chip this uses. The ESP32C3 is a 32-bit RISC-V microprocessor, which includes a whole host of I/O peripherals in addition to Wi-Fi and Bluetooth 5 it has cryptographic hardware accelerators that support AES-128/256, Hash, RSA, HMAC, digital signature, secure boot and has a hardware random number generator. And how much is it if you purchase just one? $4.99. Come on.

Anyway, that's not why I'm telling everyone this. I'm telling everyone this because a month or two ago something in one of their mailings brought me up short because it was similarly stunning and I thought you guys all needed to at least know about it. It got away from me when I went back to try to find it. But then their most recent mailing mentioned it again, so this time I'm not letting it slip past. Get a load of this:

It's called the LinkStar-H68K-1432 multi-media router. It has Wi-Fi 6, 4GB RAM & 32GB of eMMC flash storage with an SD card slot for more. It's powered by a quad-core 64-bit Cortex-A55 chip, an ARM G52 2EE GPU. There's a GPU because it can output HDMI video at 4K x 60fps. I has a USB3 port and two USB2 ports and a USB type-C that can be attached to a SATA3 drive. On the router side, aside from its dual-band 1200Mbps Wi-Fi 6, it also has 4 Ethernet ports, two running at up to 2.5Gig and another two at 1Gig. It comes with Android 11 pre-installed but also supports Ubuntu, Debian, Armbian, OpenWRT & Buildroot (which is used to build embedded Linux systems).

So what will this little pocket-sized fanless Wi-Fi 6 4-Ethernet interface router set you back? Would you believe $119. That's what got my attention. The little NetGate SG-1100 router that I love, use and have recommended is $189, and it only has three Ethernet interfaces and no Wi-Fi. This thing has 4 separate interfaces, Wi-Fi 6 and a ton more. The fact that you can drop OpenWRT onto it and have an operating, state of the art router, becomes compelling.

I want to be clear that I don't own one, I don't have time to own one, and that I'm not vouching for it in any way. Unlike the ZimaBoard, which I was happy to vouch for since I had several and I loved them, you're on your own with this thing if you should decide to take the plunge. For the right hardware tinkerer this could be so much fun... and not very inexpensive. I have the link in the show notes AND it's episode 900's GRC shortcut: https://grc.sc/900

https://www.seeedstudio.com/LinkStar-H68K-1432-p-5501.html

And BEWARE!! – If you enjoy tinkering with hardware you could easily get lost within Seeed Studio's astonishingly inexpensive array of intriguing hardware goodies.

Elon really said this

One last piece of lunacy: When asked during a scheduled Twitter Space chat this past Sunday why he bought Twitter, Elon explained his decision as follows (and I'm not making this up):

> *"I can't exactly say why, because it's one of those things where, it's like: my biological neural nets said, 'It is important to buy Twitter' and just like with a digital neural net, you can't really exactly explain why the neural net is able to understand an image or text – the collective result of the neural net says this is an important decision, or this is the right action, and my biological neural net concluded that it was important to buy Twitter, and that if Twitter was not bought and steered in a good direction, it would be a danger for the future of civilization, and so... that's why I bought it."*

Closing The Loop

Steve Gibson / @SGgrc

> *To =ALL= Security Now Listeners:*
> *I'm currently listening to Alex Stamos on Wednesday's "This Week in Google." Alex has not let anyone get a word in edgewise because he has SO MUCH amazing information to share. Without reservation, I RECOMMEND listening to this. It's FANTASTIC!*

That tweet received about three times as many likes as my weekly Security Now! show notes posting as well as 13 replies and 12 retweets. Alex was really amazingly wonderful!

SpinRite

To say that things are going well with SpinRite's alpha release testing would be an understatement considering how poorly things could easily have gone. I'm still somewhat in shock that we're very close to having a final release. I have things to fix but nothing major so far. It's really looking good. There was one posting to the newsgroups last week that I wanted to share because it makes a point that I want to drive home and it's part of the reason why I'm fired up about SpinRite's potential long term future. This person posted:

> *I have a ThinkPad Helix and the SSD is a Samsung EVO 1TB mSATA. When the SpinRite pre-release starts, it estimates 31.7 minutes for processing. However a level 2 pass, with no errors detected, takes 2:56:06 [so just shy of 3 hours] so that's more than 5 times longer [it's actually nearly 6 times longer] than the estimated time. Is that normal?*

I replied: We've found that whereas the fronts of spinning hard drives tend to be the fastest regions, because they contain more sectors around their longer outer tracks, the fronts of MANY SSDs are conspicuously slow. We posit that this is due to the presence of much more on-the-fly error correction. We first saw this using the ReadSpeed benchmark tool. One of my future plans is to locate these slow-to-read spots and selectively re-write them to restore their speed by eliminating this unseen error correction which results in a significant reduction in SSD performance.

If you **do** discover that the front of that drive is quite slow, you **could** identify the slow region and run level 3 over just that region... and it might very well speed it back up. And that would also likely increase its reliability by solidifying those sectors which might be on the verge of transitioning from very slow to unreadable.

He replied:

> *I just did a Level 3 on the entire drive with Alpha-4 and now SpinRite estimates the 1TB drive Will require 29.7 minutes, and a full Level 2 scan completed in 29:50 (which is almost six times faster than it scanned a couple of days ago). Thank you so much for creating 6.1!*

Presumably the slow down is not uniform across the entire drive. But, interestingly it's the most used front of SSDs where we're seeing the greatest performance drop. And since that's where the OS and file system live, improving their read speed should create a useful improvement. How many times have we heard that an SSD-based machine doesn't seem to be as fast as it was when it was new? That wouldn't seem logical, but this might be what's going on.

One of the things I have in store for SpinRite once v6.1 is launched and we start work on v7 is to profile the performance of mass storage media to locate and selectively repair sluggish spots. But what anyone can do – or will be able to do soon with v6.1 – is give such drives a single level 3 pass, as this person who posted did, which might significantly improve both the system's performance and its reliability by re-writing the drive's sectors to recharge their leaking storage cells.

LastPass Again

Like many other LastPass users, last week I received another note from LastPass. The note is short so I'll read it into the record:

> *Dear valued customer,*
>
> *In keeping with our commitment to transparency, we wanted to inform you of a security incident that our team is currently investigating.*
>
> *We recently detected unusual activity within a third-party cloud storage service, which is currently shared by both LastPass and its affiliate, GoTo. We immediately launched an investigation, engaged Mandiant, a leading security firm, and alerted law enforcement.*
>
> *We have determined that an unauthorized party, using information obtained in the August 2022 incident, was able to gain access to certain elements of our customers' information. Our customers' passwords remain safely encrypted due to LastPass's Zero Knowledge architecture.*
>
> *We are working diligently to understand the scope of the incident and identify what specific information has been accessed. As part of our efforts, we continue to deploy enhanced security measures and monitoring capabilities across our infrastructure to help detect and prevent further threat actor activity. In the meantime, we can confirm that LastPass products and services remain fully functional. As always, we recommend that you follow our best practices around the setup and configuration of LastPass, which can be found here.*
>
> *As is our practice, we will continue to provide updates as we learn more. Please visit the LastPass blog for the latest information related to the incident:* https://blog.lastpass.com/2022/11/notice-of-recent-security-incident/.
>
> *We thank you for your patience while we work through our investigation.*
>
> *Sincerely,*
> *The Team at LastPass*

So, there's no follow-up yet. Last time we waited exactly three weeks from the first announcement, released on August 25th until we received an update on September 15th.

The two most relevant pieces of information in this first disclosure are:

- *We have determined that an unauthorized party, using information obtained in the August 2022 incident, was able to gain access to certain elements of our customers' information.*
- *Our customers' passwords remain safely encrypted due to LastPass's Zero Knowledge architecture.*

And, as usual from all such partial disclosures, we're left wanting something more.

We don't know which “certain elements” of their customer information was inadvertently made available — but they apparently know. We know from the follow-up note from the first intrusion that the bad guys were rummaging around in some developer network that was not connected to their production systems. But now it appears that those who were doing the rummaging managed to get sufficient information, whether in the form of source code that they disclosed, or perhaps in the form of credentials which were used to perpetrate this latest breach and information exfiltration.

If LastPass lost control of their customers' billing data — names, credit cards, street addresses and so on — that would not be good. But at this point we're just speculating. Presumably in another two weeks or so we'll be told more.

Last week, after this happened, I popped on for the first 15 minutes of Tech News Weekly with Jason and Mikah to talk about this latest breach. I made the same point that I always make, which is that none of the passwords and other secret data that's being stored by any of the many competing password managers should ever be vulnerable to any breach of the data that's being stored on our behalf in the cloud. Thanks to what we once called “PIE” for Pre-Internet Encryption, which the industry now calls end-to-end encryption — though that term is become less useful as non-end-to-end systems abuse it — we're really just using a password manager's cloud service to keep our various devices synchronized.

The threats we face to our stored secrets are only on our end. In order to do its work, at some point any password manager must have at least the user's username and password decrypted for the site being visited. I don't know whether the entire password archive is decrypted as a whole or whether sites can be decrypted individually (which would seem safer). But either way, at some moment in time the data must exist in the clear in the user's browser. Way back at the start of this podcast we noted the inherent impossibility of protecting encrypted DVD video content because the player itself needed to be able to decrypt the DVD in order to play it for its owner, and the DVD's publisher ultimately had no control over the DVD player.

So, if the password manager's browser add-on were to be adulterated in some way to break its security design, or if something was able to somehow intercept its operation in the client, that could prove devastating. But it's difficult to see how a breach of the LastPass or any other password manager's cloud synching facility could ever endanger a user's always-encrypted secrets.

Of course, none of this prevents reputational damage to LastPass, and most users will have no idea what it means for all of their data to always be encrypted before it leaves their browser. They don't see anything leaving their web browser. They have no concept of the cloud or of client-side encryption. They just know that they are using this or that password manager, that this or that password manager suffered a breach, and that the press is now able to say that this is the second breach in a little over four months.

If we assume that the decision to change password managers is unwarranted — and I'm not suggesting whether it is or not — then one huge advantage any password manager has is

inertia. It's much easier to change search engines than it is to change password managers. Certainly it can be done. And the password managers have provided means to lower the bar to doing so, offering various importers of others' password archives. But it seems most likely that until users learn that someone's passwords were actually stolen, inertia will reign. The statement "*Our customers' passwords remain safely encrypted*" matters a lot, especially when changing password managers is a pain.

The listeners of this podcast are as knowledgeable and sophisticated as any, anywhere. So their decision will be fully informed. I don't know what that means for them. But I'm continuing to remain with LastPass because I have no interest in punishing them for making a mistake, and there is no indication that the security I actually require from them has ever been endangered.

That said, I'll be interested to learn more about this most recent trouble and I'll certainly share everything I learn.