Security Now! #921 - 05-02-23

OSB OMG and other news!

This week on Security Now!

This week, because the UK's Online Safety Bill continues to stir up a hornet's nest of worries and concerns within many industries, we're going to examine WhatsApp's reaction to Signal's "we plan to walk" position and Wikipedia's concerns over the Bill's age verification requirements. And, undaunted, I have another idea that might be useful! We also have a new UDP reflection attack vector, a welcome (and late) update to Google Authenticator, more NSO Group client news, a Russian OS?, the unintended consequences of releasing updates for routers that won't actually be updated, a smart move by Intel with pre-release security auditing, yet another side-channel attack on Intel CPUs, cURL's maintainer implores Windows users not to delete it, and VirusTotal gets AI.

"Security" should never be added as an afterthought

Security News

The Encryption Debate

Before we get into everything that's happened since last week's podcast, I wanted to follow-up on last week's topic about the clear collision of encryption ideologies we're in the midst of witnessing. While I don't have anything completely novel to say, I want to go on record: I adamantly hope that not one of the encryption providers backs down from their absolutist position — which I believe is the only tenable position for our industry to take.

No one wants to provide cover to any community of law breakers of any kind — certainly not anyone trafficking in child pornography or terrorism. But the fact is, that while it's not zero, the illegal use of technology represents an infinitesimal minority of the technology's total user base. Everyone else, who are law abiding users, will obtain clear benefits from access to technology which protects their privacy to the maximum degree possible.

For the past several weeks we've been talking about the thriving marketplace for commercial smartphone spyware — a market created by the very same governments who want to expand their ability to monitor, not just targeted individuals, but everyone's private communications. Sadly, government bureaucracies are too large and unaccountable to be trusted. Edward Snowden provided a wake up call and the revelations have never stopped since.

This podcast draws lessons from events, which is much more useful and enduring than a dry recitation of the weekly news. One such lesson we have seen demonstrated time and again is that if it is possible for privacy to be breached, it will be breached. We must not willingly and knowingly provide deliberately breachable tools to unaccountable governments.

The final piece of this argument is that it would not work anyway. If fully private encryption is outlawed, only outlaws will be using fully private encryption. As we've often observed here, now that such encryption already exists it's never **not** going to exist again. So the only thing that will happen if the tools everyone uses should lose their privacy by law, is that privacy will be obtained outside of the law.

Okay? So what's going to happen?

The fact that no one knows, is what makes this so interesting. There doesn't appear to be any possible way to compromise: Either governments have no officially sanctioned way to monitor communications, or they do. It's glaringly binary. If the UK's current proposal were to be enacted into law, there would presumably be some length of time provided for encrypted service providers to come into compliance. So there would appear to be three choices:

1. A company could choose to tough it out and call the UK's bluff by simply ignoring the law and continuing to offer fully encrypted and unmonitored communications.
2. Or, they decide to comply, and during the grace period they add the technology for side-channel monitoring to their product. This complies with the UK's requirement by copying all communications to a central repository for content screening. It's still unclear how the requirement to prevent "grooming" is accommodated but it begins with monitoring.
3. Or, they just say no to compromising their users' privacy. They choose to boycott the UK, removing access to their service from all UK users.

Option one, toughing it out and calling the UK's bluff doesn't appear to be practical. A piece of related news I'll share in a moment from the BBC contains this line: *"If a service does not comply with the bill, there can be serious consequences potentially including large fines, criminal sanctions for senior staff, or restricting access to a service in the UK."*

JD Supra, a legal news site, had this to say about the Online Safety Bill's penalty provisions:

> *Penalties available include fines of up to £18m or, if higher, 10% of global turnover and which for the larger providers could be significant sums. In addition it can impose business interruption measures including, ultimately, service restrictions.*
>
> *A particularly controversial measure has been the availability of up to two years imprisonment for senior managers who suppress, destroy or alter information requested by OFCOM, who fail to comply with, obstruct or delay OFCOM when exercising its powers of entry, audit and inspection, for providing false information or for employees who fail to attend or provide false information at an interview. A recent amendment also provides a further offense where a senior manager has "consented or connived in ignoring enforceable requirements, risking serious harm to children".*
>
> *For these purposes a "senior manager" is if the individual plays a "significant role in (a) the making of decisions about how the entity's relevant activities are to be managed or organized, or (b) the actual managing or organizing of the entity's relevant activities."*
>
> *The OSB is in the latter stages of the legislative process and while substantive amendments may still be made, it is likely to receive Royal Assent by mid-year. The OSB may also have the unintended effect of causing terms of service to be watered down as to what content a service may contain. Ultimately, some providers may decide it is simply too difficult to comply with and instead block UK users.*

There's obviously no way Tim Cook or Mark Zuckerberg are going to prison over this, nor are they going to give up 10% of their respective companys' annual revenue for the privilege of offering their UK users un-monitored communications. So ignoring the law doesn't appear to actually be a practical option, and the law was clearly crafted to make exactly that fact exceedingly clear. But I wanted to cover that case, since it was theoretically one of the three possibilities.

Age does matter...

It turns out that I wasn't totally out in the weeds last week with my proposed design for a device-centric age verification solution; and something like that may actually be the right answer. Even if the Online Safety Bill fails, as we all hope it will, in forcing the monitoring of all private communications within the UK—which, regardless of how the Bill's authors want it to be "interpreted", is what the Bill requires—the Bill's extensive provisions, which require the moderation of all content made available to those under the age of 18, appear very likely to survive. So it appears that some answer to the unsolved challenge of online age verification will be needed.

This was highlighted by some news that was covered four days ago by the BBC in their piece titled: *"Wikipedia will **not** perform Online Safety Bill age checks."*

Here's what the BBC wrote:

The Wikipedia foundation says that it will not comply with any age checks required under the Online Safety Bill. Rebecca MacKinnon, of the Wikimedia Foundation, which supports the website, says it would "violate our commitment to collect minimal data about readers and contributors".

A senior figure in Wikimedia UK fears the site could be blocked as a result. But the government says only services posing the highest risk to children will need age verification. Wikipedia has millions of articles in hundreds of languages, written and edited entirely by thousands of volunteers around the world.

It is the eighth most-visited site in the UK, according to data from analytics company SimilarWeb. The Online Safety Bill, currently before Parliament, places duties on tech firms to protect users from harmful or illegal content and is expected to come fully into force some time in 2024.

Neil Brown, a solicitor specializing in internet and telecom law, says that under the bill, services likely to be accessed by children must have "proportionate systems and processes" designed to prevent them from encountering harmful content. That could include age verification.

Lucy Crompton-Reid, chief executive of Wikimedia UK, an independent charity affiliated with the foundation, warns some material on the site could trigger age verification. "For example", she said, "educational text and images about sexuality could be misinterpreted as pornography."

But Ms MacKinnon wrote: "The Wikimedia Foundation will not be verifying the age of UK readers or contributors." As well as requiring Wikipedia to gather data about its users, checking ages would also require a "drastic overhaul" to technical systems. If a service does not comply with the bill, there can be serious consequences potentially including large fines, criminal sanctions for senior staff, or restricting access to a service in the UK. Wikimedia UK fears that the site could be blocked because of the Bill, and the risk that it will mandate age checks.

Ms Crompton-Reid wrote: It is "definitely possible that one of the most visited websites in the world - and a vital source of freely accessible knowledge and information for millions of people - won't be accessible to UK readers (let alone UK-based contributors)." There are currently 6.6 million articles on Wikipedia, and she said it was "impossible to imagine" how it would cope with checking content to comply with the bill.She added: "Worldwide there are two edits per second across Wikipedia's 300-plus languages."

The foundation has previously said the bill would fundamentally change the way the site operated by forcing it to moderate articles rather than volunteers. It wants the law to follow the EU Digital Services Act, which differentiates between centralized content moderation carried out by employees and the Wikipedia-style model by community volunteers. Last Tuesday, the House of Lords debated an amendment from Conservative peer Lord Moylan that would exempt from the Online Safety Bill services "provided for the public benefit," such as encyclopedias. Heritage Minister Lord Parkinson said he did not think this would be feasible,

but added that Wikipedia was an example of how community moderation can be effective. He said the bill did not say that every service needed to have age checks, and it was expected that "only services which pose the highest risk to children will be required to use age verification technologies."

Ms Crompton-Reid told the BBC that while Lord Parkinson's remarks "reassured" her, the charity did not want to be relying on future goodwill and interpretation of legislation. She said they would continue to urge that protections to community moderation were in the bill through measures such as an exception for public benefit websites like Wikipedia.

So our takeaway from that is that while Wikipedia will refuse to comply with any age-related regulation, it seems unlikely that Wikipedia would actually be required to do so—it's not a porn site. They would appreciate receiving confirmation that this will not be necessary in order to remove any uncertainty. But, there certainly are services that will be required to provide age verification, such as legal websites which make it their business to provide extremely adult sexual content.

Age Verification

Last month, The Guardian published what they called an "Explainer" about the age-related aspects of the impending Online Safety Bill. Their piece was titled *"Will UK's online safety bill protect children from adult material?"* with the sub-heading: *"Legislation puts duty of care on tech firms to protect under-18s but does not mandate use of specific age-checking technology."* They wrote:

The online safety bill is due to become law this year and it imposes a duty of care on tech companies to protect children from harmful content. However, there are calls from campaigners and peers to toughen the legislation's provisions regarding pornography. Here is what the act proposes to do on adult material.

The bill requires all pornography websites, such as Pornhub, to ensure children do not encounter their content. This will require age-checking measures. The legislation refers to stringent age verification – checking a user's age via government ID or an authoritative data source such as a person's bank – as a means of doing so. Breaches of the act carry the threat of a fine of up to 10% of a company's global turnover or, in extreme cases, blocking a website altogether.

So, what are the rules currently?

MPs have described the legal approach to pornography in the UK as a "loose patchwork" comprising more than a dozen laws. It is a criminal offense to publish work under the Obscene Publications Act that is deemed "obscene" and it is illegal under the Criminal Justice and Immigration Act to possess an "extreme" pornographic image. It is also an offense to make, possess or distribute indecent images of a child.

The primary regulator of legal pornography ***offline*** *is the British Board of Film Classification, which gives pornography age ratings – R18 for the most extreme but legal content, or 18 – but it has no control over online content.*

Ofcom, the communications watchdog, already has the power to regulate UK-based "video-sharing platforms" such as TikTok, Snapchat and OnlyFans. These platforms are required to protect under-18's from videos containing R18 material such as pornography.

The age appropriate design code was introduced in 2021 and is designed to prevent websites and apps from misusing children's data. Under its terms, social media platforms would be breaching the code if their algorithms served adult material to under 18-year-olds.
Age verification has been a troublesome issue for the government.

[And I'll just note here that "online age" is an aspect of "online identity" and online identity is both a thorny problem and a slippery slope that we haven't even begun to solve.]

Age-checking for pornography was announced as a Conservative policy in 2015. However, plans to introduce a nationwide age verification system for online pornography were abandoned in 2019.

The bill will not mandate use of specific technologies for age checking, although Ofcom will issue codes of practice on age assurance, which is the umbrella term for assessing the age of people online. Age verification is the term for the toughest measures, such as requiring proof of official ID.

One solution is to use age verification companies that vet a user's age – via a range of methods including checking official ID or bank statements – and then notify the porn provider that the person wishing to access their service, who is anonymised, is over 18 years old.

Ofcom has said it will launch a consultation on protecting children from pornographic content – including on user-generated platforms such as OnlyFans – in the autumn.

The government has indicated that there will be clear instructions to mainstream social media sites and search engines to prevent children accessing pornographic content on their services. The bill requires sites to prevent children encountering what it terms "primary priority content". Because it qualifies as a "user-to-user" service, subscription site OnlyFans is also covered by this part of the bill.

We will not know what is primary priority content officially until it is defined in a statutory instrument that will be published after the bill becomes law. However, pornography is expected to be on that list and it was listed as primary priority content by the previous culture secretary, Nadine Dorries, in a parliamentary statement last year. According to a timeline published by Ofcom, though, it could be more than 18 months after the bill is passed before these provisions come into effect.

Social media sites and legal pornography sites will also be required to shield all users from illegal pornography such as obscene content and child sexual abuse material.

The bill will update the law on sharing intimate images without someone's consent. In England and Wales there will be a new "base offense", where it is an offense to share an intimate image of a person if they do not consent – and the perpetrator does not believe they have consented. Currently, these offenses apply if the image is shared in order to cause humiliation or distress.

The base offense will now apply regardless of the motivation, including sharing it as a joke, for social status, financial gain or "where there is no motivation at all".

From all of that, the article's sub-heading seems the most pertinent to those of us who are interested in how things are done. The Guardian wrote: *"Legislation puts duty of care on tech firms to protect under-18s, but does not mandate use of specific age-checking technology."* Saying that another way... *"We don't know how you're going to arrange to do, what the new laws we have just written require you to do. But that's not our problem, it's yours, because we said so."* Whether it's the perceived need to monitor everyone's communications all the time in the off chance that something illegal might pass by, or the need to impose strict age restrictions on access to Internet content and behavior, it's apparent that the UK's legislators believe that they can ask for whatever they want, leaving it up to the tech companies to figure out how to do it, all the while imposing penalties if they fail to achieve what might well be impossible.

After last week's podcast, I received a DM from someone who said: *"**Please** don't go spending another seven years solving the age verification problem like you did with the online login authentication problem."* To which I will formally respond: "Fear not. As they say, 'fool me once' I have quite thoroughly learned my lesson. And I'm having far too much fun working on new technology for SpinRite, where I can actually make a difference."

WhatsApp: Rather be blocked in UK than weaken security

Recall that Signal's president Meredith Whittaker made some headlines when she told BBC News that signal *"would absolutely, 100% walk"* and stop providing services in the UK if required, by the Online Safety Bill, to weaken the privacy of its encrypted messaging system.

That's the only stance that any entity like Signal, Threema or Telegram could take. But what's the position of the number one most popular and largest messaging app in the UK, WhatsApp? WhatsApp is used by more than seven in 10 adults who are online, according to the UK's communication regulator Ofcom.

So, the BBC asked Will Cathcart, the head of WhatsApp. Will replied that **WhatsApp would refuse to comply if asked to weaken the privacy of encrypted messages.** Period. Full stop. He said that WhatsApp would rather be blocked in the UK, than undermine its encrypted-messaging system, if required to do so under the Online Safety Bill. He said: *"We won't lower the security of WhatsApp. We have never done that - and we have accepted being blocked in other parts of the world."* And he feared the UK would set an example other nations might follow. Will added that undermining the privacy of WhatsApp messages in the UK would do so for all users: *"Our users all around the world want security - 98% of our users are outside the UK, they do not want us to lower the security of the product. We've recently been blocked in Iran, for example. We've never seen a liberal democracy do that. When a liberal democracy asks, 'Is it OK to scan everyone's private communication for illegal content?' that emboldens countries around the world that have very different definitions of illegal content to propose the same thing."*

After Will Cathcart went on the record with WhatsApp's position, Signal's Meredith Whittaker tweeted: *"looking forward to working with @wcathcart and others to push back"*, after which Will replied on Twitter: *"And very important we work together (and honoured to get to do so) to push back."* So, what appears to be forming is a bit of an insurrection and this may be where the encrypted services companies decide they need to take a stand. Last week we saw their open

letter to the UK regulators, so we know they all know each other, they have each other's eMail addresses, and they're talking. That's all for the good.

I wondered what Apple might do, since iOS's always-encrypted iMessage is so deeply integrated into their their products. Then I considered the green bubbles. Assuming that Apple also decides to "just say no" to government communications monitoring, they could simply drop the use of iMessage encryption and fall back to SMS whenever they're communicating inside the UK. So it would mimic the way iOS devices currently operate when messaging outside of Apple's closed and encrypted ecosystem to Android devices.

Exposing Side-Channel Monitoring

I want to get off this topic now, but first there's one more thing I need to share. It's just a concept and observation that I want to plant in everyone's mind. Given WhatsApp's stance, which aligns with all of the other encryption providers, I doubt this idea will be needed. I truly hope that's the case, and I believe that if everyone **just says no**, that's likely to work. I doubt that the citizens of the United Kingdom would choose to be without all of their messaging capabilities with each other and with the rest of the world, especially when having that happen would only serve to drive the creeps further underground. But, if just saying "no" doesn't work for some reason, we may need a fallback.

One of the mixed blessings of today's technologies is that most people have no idea how they operate. And for the most part that's good. You should not need to be a car mechanic to drive a car. That's the leverage provided by technology. But this also means that most car drivers have very little idea what's going on under the hood. If they don't need to know, then not knowing is fine. But if there's a problem, some knowledge could come in handy when it comes to making decisions.

For quite a long time, third-party cookies lived in obscurity. They were always there, but by design they remained part of what was under the hood, out of sight and out of mind. And third party cookies liked it that way. But no one who was asked whether they wanted to have third-party tracking cookies said that they thought that would be a great idea.

Tracking has a similar history. For years it has been going on largely unseen (often aided by those same third-party cookies), and it has enabled an entire online web surveillance industry. But when Apple began requiring that iOS applications obtain explicit permission to allow the app's users to be tracking outside of the application, the result was an overwhelming and resounding cry of NO THANK YOU! (And that's putting it politely.)

And I had my own firsthand experience with people being unhappily surprised. Leo used to introduce me by mentioning that I discovered the first spyware and in the process coined that term. What happened was that I discovered that a freeware utility I had installed on my Windows machine – as I recall it was an early version of WinZIP – wasn't as free as I was led to believe. It was an "ad supported application", and so it brought along an advertising DLL from a third-party company named "Aureate." What I discovered was that this Aureate spyware was inventorying my machine, monitoring my actions and usage, and then phoning home without ever obtaining any permission from me. I had no idea it was there, and to say that I was

unhappy when I found something communicating behind my back without my knowledge or permission, would be an understatement. So I created "OptOut", the world's first spyware removal tool, a bit of freeware which successfully removed Aureate and several other early forms of stealth spyware. The reason I bring this up is that the management at Aureate shared with me some of the way over-the-top enraged and nearly psychotic eMails they were receiving from PC users who were more than just a bit unhappy to have used OptOut to discover that **their** machines had also been infected. The Aurate people said that the ad-supported software packages which installed their spyware – I mean their ad-ware – were supposed to explain the situation and obtain their users' consent. I asked why they didn't have their DLL present its own permission dialog. Anyway, the name "Aurate" had been ruined by my crusade, so the company changed its name to Radiate, and not long after ceased operations. The whole concept was really never viable and no freeware developer wanted anything to do with them.

So what does all this have to do with the UK's Online Safety Bill?

What occurred to me was that if encrypted applications were going to be required by law to arrange some sort of side-channel government mandated monitoring and surveillance, they should make very clear to their user that that's what they are doing and not repeat past mistakes of doing things that people would find objectionable if they were clearly informed of what was going on. So, the presence of state-mandated communications surveillance should be placed front and center for every UK resident – and anyone they communicate with. The top of any application that's being forced to break into its user's privacy in order to comply with the UK's Online Safety Bill should clearly display against a red background, the message:

This communication is being monitored by your government

I imagine that the presence of that notice at the top of any communicating application might provoke a reaction similar to what happened when the news of the Aureate spyware broke. The UK government will clearly wish the fact of this be hidden from its citizen's view. But it should be there to serve as a constant reminder of what the country's politicians have decided is in the best interests of their citizenry.

Closing the Loop

tws / @twshaka

Hi Steve,
Are you still using the ZimaBoard you mentioned many episodes ago? Would you still recommend? With the shortages of pi's I'm considering them? Thanks for making spin right and the SN podcast.

In a word, Yes yes yes! Oh wait, that's three words. The ZimaBoard is the best thing I've found for my own work during the development of SpinRite. It's perfect. And while I haven't taken a show of hands in the SpinRite dot Dev newsgroup, I keep seeing people referring to their ZimaBoard in passing. So I know that I'm not alone. But also for other uses, it really is a perfect

little machine. It's been around long enough that there are now a ton of YouTube How-To videos covering pretty much anything you can think of. And, while working to assemble today's show this morning, I received an eMail from them announcing a Star Wars day discount of 20%. Their eMail said: "May the 4th be with you" (apparently there's a little bit of a listhp) and their eMail says that the sale runs for three days, from tomorrow, May 3rd through Friday the 5th, offering, again, a 20% discount. So, yeah… ZimaBoard dot com. And if I didn't already own five of them, I'd be purchasing more. It was a real find.

Remember, though, that unlike the Raspberry Pi, which is ARM-based, the ZimaBoard is Intel based. That's what makes it incredibly useful to me... but you'll want to be sure that whatever you want to do with it you can do with Intel-based software. :) And definitely check out YouTube!

David Schofield / @hotspotoffice

Good morning Steve! SN listener since 2007, Spinrite user since my ST225's! 😁 Encouraging news about Microsoft rewriting parts of windows in memory-safe Rust!

This is a perfect segue for mentioning a bit of news that I had encountered. David's DM linked to an article in "The Register." Here's a bit from their piece:

Microsoft is rewriting core Windows libraries in the Rust programming language, and the more memory-safe code is already reaching developers. David Weston, director of OS security for Windows, announced the arrival of Rust in the operating system's kernel at BlueHat IL 2023 in Tel Aviv, Israel, last month. He said: "You will actually have Windows booting with Rust in the kernel in probably the next several weeks or months, which is really cool. The basic goal here was to convert some of these internal C++ data types into their Rust equivalents."

Microsoft showed interest in Rust several years ago as a way to catch and squash memory safety bugs before the code lands in the hands of users; these kinds of bugs were at the heart of about 70 percent of the CVE-listed security vulnerabilities patched by the Windows maker in its own products since 2006. The Rust toolchain strives to prevent code from being built and shipped that is exploitable, which in an ideal world reduces opportunities for miscreants to attack weaknesses in software. Simply put, Rust is focused on memory safety and similar protections, which cuts down on the number of bad bugs in the resulting code.

Rivals like Google have already publicly declared their affinity for Rust. Amid growing industry support for memory safe programming, Microsoft's exploration of Rust has become more enthusiastic. And last September, it became an informal mandate: Microsoft Azure CTO Mark Russinovich declared that new software projects should use Rust rather than C/C++.

All of the evidence suggests that we're not really making any headway with simply trying to be more careful using powerful but unsafe legacy languages. Our software is growing more and more complex, and people make mistakes. The adoption of newer languages which prevent those mistakes from proving fatal to a system's security looks like the only way we're going to start removing **more** existing bugs than new bugs we introduce.

Frank S / @frank_f1w

Dear Steve, I have two kids and I am very happy that they are under 7. This gives me and the market some time to find best (or better) practices for children using the internet.

As a parent, I want my children to be safe. Both physically and online. However, I don't think that we can stop existing CSAM images that are already out there. What we can and should do, is try to prevent new cases and victims.

I believe it is up to the parents to take better care and guide our children when growing up. And I don't want the government to spy on them. As a parent I would like more tools to keep my children safe online.

I wanted to use this as a catalyst to thank all of our listeners who took the time to write after last week's podcast. It's clear that everyone understands that the Internet is a true mixed blessing when it comes to our youth who haven't yet obtained the life experience which would allow them to place the horrific crap they might encounter on the Internet into its proper context. Unfortunately, bad and taboo things can be exciting and excitement can be addictive.

I was brought up short by a tweet I received which noted that, unfortunately, giving parents control and oversight over their children's communications could be harmful to the child when the child's nature is rejected and not understood by their parents. In such situations, having private communications creates a sanctuary. I think my answer to that situation, which I can certainly imagine and empathize with, is to observe that the Internet didn't create such problems, it's just another part of a complex world. And I think that it's easy to make the mistake — which I would argue UK legislators are making — of assuming that all such problems can be solved by the proper application of technology. I'm certain that's not true, and I think it's possible to get ourselves tied up in knots trying to whack every mole.

We've been focused upon the UK so far, but this is probably a good time to mention that a new US Federal bill was announced and unveiled last Wednesday aiming to regulate access by age, to social media platforms in the US. Here's a bit of CNN's coverage of the this new proposal:

A new federal bill unveiled Wednesday would establish a national minimum age for social media use and require tech companies to get parents' consent before creating accounts for teens, reflecting a growing trend at all levels of government to restrict how Facebook, Instagram, TikTok and other platforms engage with young users.

The proposed legislation by a bipartisan group of US senators aims to address what policymakers, mental health advocates and critics of tech platforms say is a mental health crisis fueled by social media.

Under the bill, known as the Protecting Kids on Social Media Act, social media platforms would be barred from letting kids below the age of 13 create accounts or interact with other users, though children would still be permitted to view content without logging into an account, according to draft text of the legislation.

Tech platforms covered by the legislation would also have to obtain a parent or guardian's consent before creating new accounts for users under the age of 18. The companies would be

banned from using teens' personal information to target them with content or advertising, though they could still provide limited targeted recommendations to teens by relying on other contextual cues.

It's the latest step by lawmakers to develop age limitations for tech platforms after similar bills became law this year in states such as Arkansas and Utah. But the legislation could also trigger a broader debate, and possible future court challenges, raising questions about the privacy and constitutional rights of young Americans.

Speaking to reporters Wednesday, Hawaii Democratic Sen. Brian Schatz, an architect of the federal bill, said Congress urgently needs to protect kids from social media harms. Schatz said: "Social media companies have stumbled onto a stubborn, devastating fact. The way to get kids to linger on the platforms and to maximize profit is to upset them — to make them outraged, to make them agitated, to make them scared, to make them vulnerable, to make them feel helpless, anxious [and] despondent."

And I'll just note that this discovery is not unique to social media platforms appealing to children. Precisely the same observation has been made by cable news outlets about how to engage and enrage their audiences to encourage viewership. I'm not sure what "gen" we're on now, X, Y, or Z... but young people have grown up with the Internet and hundreds of cable channels each with their own agenda. I hope they can figure out how to handle the mess we've made of this.

Speaking of messes...

Simon Zerafa 🖖 / (@simonzerafa@infosec.exchange) @SimonZerafa

Shodan: "Twitter canceled our API access which broke the ability to login to Shodan via Twitter (SSO) . Email us at support@shodan.io if you're currently logging in via Twitter and would like to migrate to a regular Shodan account instead of using SSO."

Yikes! This reminds us that aside from the well appreciated privacy implications of using OAuth-style "Sign in with..." authentication. If that third-party authentication service should ever become unavailable for any reason, you're kind of hosed since the only way you're known by the service you're wishing to login to is courtesy of that other entity which may no longer exist.

The Week's News

A new UDP reflection attack vector
We've previously talked about UDP reflection attacks. Unlike TCP connections, which are inherently bi-directional and therefore require packet round trips between the endpoints to establish byte numbering and other connection parameters, the UDP protocol is often referred to as a "connectionless" protocol because, although it's still possible to establish connections by mutual agreement, UDP doesn't have that baked in. This makes UDP the perfect protocol for DDoS bandwidth flooding attacks since the sender of a UDP query can spoof the UDP packet's own source IP address so that the recipient of that query will direct its reply to the victim of the attack.

What's then needed are publicly exposed and available UDP protocol services which will generate a large answering reply to a very small query. This is known as the UDP query amplification factor. How many times larger is the reply than the query. A not very exciting example of such a service is good old DNS. A relatively small query for a DNS record can return a significantly larger answer. But DNS is optimized for small size and its internal compression is really quite clever. So, as I said, it's not very exciting.

We're revisiting this subject because security researchers from Bitsight and Curesec have stumbled upon a way to exploit a network service that was only ever intended for internal LAN use, but for which, for some reason, about 70,000 instances are currently exposed on the public Internet. The service in question is the Service Location Protocol (SLP) and by cleverly abusing it through the means that these researchers discovered — and have now published with full exploit details — UDP query amplification factors as high as 2200 to 1 can be achieved. This makes this technique one of the largest amplification factors ever discovered. And since the service is available over UDP, it is ready made for DDoS flooding.

Because of the protocol's huge potential for DDoS attacks, both Cloudflare and Netscout have said that they expect the prevalence of SLP-based DDoS attacks to rise significantly in the coming weeks once threat actors learn to exploit it. The only good news is that since SLP is transported over port 427, and since it has no business being out on the public Internet, I would expect that major highly savvy carriers like CloudFlare will already be proactively blocking that traffic at their borders. But won't help any unprotected targets. So DDoSers have had another arrow added to their quiver.

Google Authenticator Updated
Google authenticator, which was first released 13 years ago in 2010, has just been updated with an extremely useful new feature which I've had in my favored iOS OTP Auth app from the start: Cloud backup. In their announcement of this groundbreaking technology, they wrote:

One major piece of feedback we've heard from users over the years was the complexity in dealing with lost or stolen devices that had Google Authenticator installed. Since one time codes in Authenticator were only stored on a single device, a loss of that device meant that users lost their ability to sign in to any service on which they'd set up 2FA using Authenticator.

With this update we're rolling out a solution to this problem, making one time codes more durable by storing them safely in users' Google Account. This change means users are better protected from lockout and that services can rely on users retaining access, increasing both convenience and security.

So, that's great. Somewhat odd that it took them this long, but I wanted to make sure that everyone listening who might be a Google Authenticator user, as I once was, would know of this critically useful new feature.

Does Israel use NSO Group commercial spyware?

I suppose we should not be surprised that Israeli law enforcement is apparently using their own NSO Group's spyware to spy on their own citizens. In response to reports in Israeli media which claim that their police have been using a "reduced strength" version of the NSO Group's Pegasus spyware—known as Saifan—to target activists, business figures, reporters, and politicians... perhaps to save face, the Israeli government has announced the formation of a commission to probe into the use of spyware by police forces to hack the smartphones of Israeli citizens. Isn't the definition of a "commission" the place where sensitive political issues go to die? The announcement makes everyone happy and no one can say that the government isn't doing anything *"Hey! There's a commission for that!"* Then, after a few years of inaction it will be quietly disbanded. In any event, if the Israeli media reports are accurate, it was interesting that Israel's own police are also getting in on the act. But the privacy of activists, business figures, reporters, and politicians is being breached.

A Russian OS?

The article in Russia's Kommersant news was written in Russian and I didn't see any way to easily translate it into English. But the news is that the Russian government is working on a law to force retailers to pre-install Russian operating systems on all new PCs sold in the country — instead of Windows. The first wave of feedback claims this will lead to an increase in laptop and PC prices across Russia. Despite efforts to get Russian companies and users to move to Russian operating systems, Windows' market share remained the same in Russia as last year. It would be interesting to see what a Russian operating system looks like. It must be a derivative of Linux. That's the only thing that would be feasible. But then, if so, why would that be more expensive than Windows?

TP-Link routers compromised

During the Toronto Pwn2Own hacking contest which we covered last December, one of the successfully exploited devices was a fully patched (at the time) TP-Link router. After the exploit was created and demonstrated during the contest event, it was assigned a CVE and the contest organizers, ZDI, the Zero Day Initiative, responsibly reported that vulnerability to TP-Link. TP-Link found and fixed the trouble and released a patch for it this past march. And, unfortunately, that patch was all the operators of the Mirai DDoS botnet needed in order to reverse engineer the change to discover the flaw for which that patch had just been released. They immediately then set about taking over and hijacking every TP-Link router that had not yet

been updated.

So we have a story where everyone did everything right. Everyone acted correctly. A problem was found. It was demonstrated. The details were kept secret and the underlying flaw was responsibly reported to the product's publisher who, in a somewhat timely manner, fixed the trouble and made an update available to their devices. But despite everyone doing everything exactly right, the bulk of TP-Link routers were almost certainly never updated and with today's Internet scanner databases, discovering the locations of those routers is not difficult.

The evidence suggests something that's obvious in retrospect: Bad guys are watching every minute of hacking contests such as Pwn2Own. They're just waiting to see someone hack something where there will be a large "patch gap" which exists between the eventual release of an update and those updates being installed into vulnerable gear. And probably nowhere is the "patch gap" larger and more glaring than in consumer routers. When was the last time any of us checked to see whether our router had new firmware available? I just checked and, sure enough, my ASUS consumer router has newer firmware available. But how would I know that? I'm not obsessively checking it every day.

In retrospect and perversely, it would almost have been better if TP-Link had not published a public update for their firmware because the act of doing so painted a big red bull's-eye on every publicly exposed vulnerable TP-Link router — and the phrase "publicly exposed" is redundant for a router, since that's what they are almost by definition. Assuming that the problem was present in their current product line, TP-Link might have simply fixed it there and never pushed out a fix for this problem.

I know that goes against everything we think and believe about fixing and updating known problems. But if patches cannot reasonably be expected to be applied then what will happen is what just did. If the unknown problem had been left alone, all of those vulnerable routers would not now be Pwned and enslaved into the Mirai botnet.

I appreciate the controversy surrounding this, but I think that generic consumer routers all need to occasionally phone home to check for updates and be willing to take themselves offline for an autonomous update cycle. The clear prevalence of bad guys who are now waiting to receive and reverse engineer router patches, coupled with the fact that router owners don't know to patch, I think tips the balance in favor of such routers being autonomously self-updating.

A pre-release security audit:
So, we have a bit of happy Intel news. First we need to know what Intel's TDX is: TDX stands for "Trusted Domain Extensions". Intel describes it as:

> *Intel Trust Domain Extensions (Intel TDX) is introducing new, architectural elements to deploy hardware-isolated, virtual machines (VMs) called trust domains (TDs). Intel TDX is designed to isolate VMs from the virtual-machine manager (VMM)/hypervisor and any other non-TD software on the platform to protect TDs from a broad range of software.*
>
> *VM isolation with Intel TDX is a key component of Intel's Confidential Computing portfolio,*

which also includes application isolation with Intel SGX and trust verification with our upcoming service code-named Project Amber. Confidential Computing uses hardware to protect data in-use from a wide variety of threats and enables organizations to activate sensitive or regulated data that may have otherwise been locked down and idle.

After this, Intel brags that **before** releasing this new tech to the world, they ran it through a very useful security gauntlet. They wrote:

- *In our first-ever pre-release activity, we also took Intel TDX through Project Circuit Breaker, part of Intel's Bug Bounty, where we challenged a community of elite hackers to find bugs in some of our top technologies. Using simulation software, the community went through two rounds of bug hunting over several months, earning bounties to help us find potential vulnerabilities so we could mitigate them*

- *We then took it to security experts at Google Cloud and Google Project Zero to conduct a deep security review. They looked for security weaknesses while evaluating the expected threat model for any limitations that would inform Google's decisions. The 9-month collaboration resulted in 10 security issues and 5 defense-in-depth changes that were mitigated.*

- *Intel offensive researchers also spent considerable time reviewing the product. Their job is to apply an attacker mindset to evaluate security technologies. They were able to find and mitigate potential vulnerabilities like the use of memory disturbance errors. Threat modeling, penetration testing, and hackathons were all applied during the research.*

So, the good news is that even though the authors of Intel's code, were as sure as any code authors ever are that their code was correct, they nevertheless subjected it to pre-release third-party scrutiny. Naturally, Intel put a happy face on the results, saying that it had succeeded in improving the code quality. What we learn from other sources is that, indeed, vulnerabilities were uncovered during the security audit that could have resulted in arbitrary code execution, cryptographic weaknesses, and denial-of-service.

Another Intel side-channel attack:
An academic paper was just published titled: *"Timing the Transient Execution: A New Side-Channel Attack on Intel CPUs"*. Since we've entered the world of "yet another side-channel information leakage from Intel chips" I'm not going to spend undue time digging into this one, but in case it might wide up being important—which seems unlikely—I wanted to at least share their brief description, which reads:

The transient execution attack is a type of attack leveraging the vulnerability of modern CPU optimization technologies. New attacks surface rapidly. The side-channel is a key part of transient execution attacks to leak data. In this work, we discover a vulnerability that the change of the EFLAGS register in transient execution may have a side effect on the Jcc (jump on condition code) instruction after it in Intel CPUs. Based on our discovery, we propose a new side-channel attack that leverages the timing of both transient execution and Jcc instructions to deliver data. This attack encodes secret data to the change of register which makes the execution time of context slightly slower, which can be measured by the attacker to decode

data. This attack does not rely on the cache system and doesn't need to reset the EFLAGS register manually to its initial state before the attack, which may make it more difficult to detect or mitigate. We implemented this side-channel on machines with Intel Core i7-6700, i7-7700, and i9-10980XE CPUs. In the first two processors, we combined it as the side-channel of the Meltdown attack, which could achieve 100% success leaking rate. We evaluate and discuss potential defenses against the attack. Our contributions include discovering security vulnerabilities in the implementation of Jcc instructions and EFLAGS register and proposing a new side-channel attack that does not rely on the cache system.

So, yeah, once again, aspects of our modern processors which were developed to improve their performance in common contexts, and which existed for years with no one worrying, are one by one turning out to be exploitable to leak information wherever hostile code might be sharing hardware with targeted code which contains secrets. And I'll note that these days virtually all operating systems contain valuable data that needs to be kept secret.

Windows users: Don't remove cURL!

Early last week, Daniel Stenberg posted two messages at Mastodon.social regarding some recent hysteria about cURL. Daniel's feelings about cURL are significant because cURL is pretty much his baby. His own bio says: "I am the founder and lead developer of cURL and libcurl. An internet protocol geek, an open source person and a developer. I've been programming for fun and profit since 1985. You'll find lots of info about my various projects on these web pages and on my GitHub profile. My name appears in products." Daniel was also the 2017 winner of Sweden's prestigious Polhem Prize for his work on cURL. Here's what Daniel posted to Mastodon

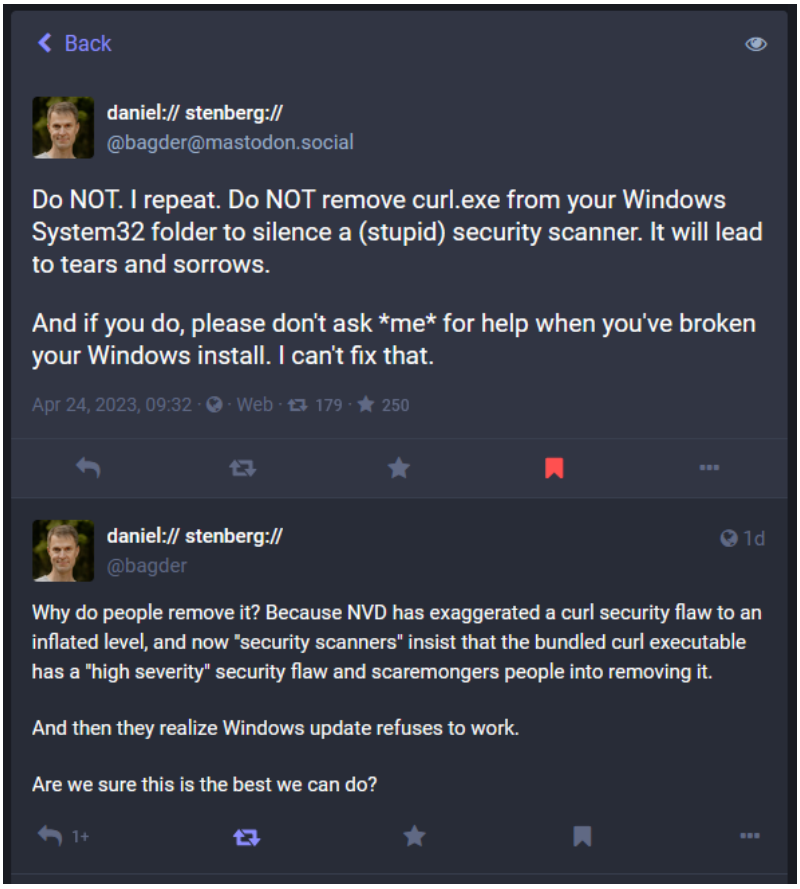

And, finally…

AI comes to VirusTotal:
During last week's annual RSA security conference, Google announced "VirusTotal Code Insight." Code Insight is a new feature for VirusTotal that uses AI/Machine Learning to generate simple natural language summaries from submitted malware and source code samples. Google says that, at present, the new functionality is deployed to analyze PowerShell files submitted to VirusTotal. The company says it plans to expand the service with additional file formats in the future.

This is potentially very cool and useful. However, increasing real world evidence suggests that AI (at least in the form of ChatGPT) always sounds absolutely convincing and authoritative while being completely wrong with its "facts." So, at least at this early stage anything and everything that's produced should be regarded with skepticism and carefully verified.

But, since this is a Google effort and it would be truly cool to have something said through an automated analysis of uploaded malware, I have hoped for the future of this.