Security Now! #920 - 04-25-23

An End-to-End Encryption Proposal

This week on Security Now!

This week's look at the past week's most interesting security news answers the question of whether Apple's Lockdown Mode does anything that's actually useful? Just how big is the market for commercial "Pegasys-style" smartphone spyware? Why exactly has the Dark Web suddenly become interested in purloined ChatGPT accounts and is "purloined" a word one uses in mixed company? What trove of secrets did ESET discover when they innocently purchased a few second hand routers? And speaking of routers, what was the mistake that users of old Cisco routers really wish Cisco hadn't made, and whose fault is its exploitation today? What's the story behind the newly established Security Research Legal Defense Fund? Then, after a few quick update and upgrade notes, we look at two opposing open letters written about the coming end-to-end-encryption apocalypse, and consider whether I may have just stumbled upon a solution to the whole mess? So, I doubt that anyone's going to be bored this week!

By some miracle, it's still working... ***DON'T TOUCH IT!***

Security News

Lockdown Mode seen succeeding

Last Tuesday, the forensic security research group Citizen Lab reported on three iOS 15 & 16 exploits attributed to Israeli's NSO Group's Pegasys smartphone spyware system.

As a reminder, "The Citizen Lab" is at University of Toronto's Munk School of Global Affairs and Public Policy. They've been doing some serious forensics work. Last week's main topic was "Forced Entry", a fascinating examination of the details of a zero-click exploit against iOS devices which Google's Project Zero researchers reverse engineered and dissected thanks to Citizen Lab finding a live sample of it on the phone of a Saudi political activist.

So, my eye was caught by Citizen Lab's mention last week of the apparent successes of Apple's Lockdown Mode which we've previously described. Lockdown Mode rather significantly restricts many features of an iPhone for the express purpose of thwarting exactly these sorts of targeted attacks against high profile users of iOS devices.

Last Tuesday, they published an extensive description of several zero-click attacks they discovered being deployed against users of iOS 15 and 16. I'm not going to get into all of the details of those, but their intersection with Lockdown Mode is interesting. They summarized their finds in seven bullet points which are short and, I think, worth sharing:

- *In 2022, the Citizen Lab gained extensive forensic visibility into new NSO Group exploit activity after finding infections among members of Mexico's civil society, including two human rights defenders from Centro PRODH, which represents victims of military abuses in Mexico.*

- *Our ensuing investigation led us to conclude that, in 2022, NSO Group customers widely deployed at least three iOS 15 and iOS 16 zero-click exploit chains against civil society targets around the world.*

- *NSO Group's third and final known 2022 iOS zero-click, which we call "**PWNYOURHOME**," was deployed against iOS 15 and iOS 16 starting in October 2022. It appears to be a novel two-step zero-click exploit, with each step targeting a different process on the iPhone. The first step targets HomeKit, and the second step targets iMessage.*

- *NSO Group's second 2022 zero-click ("**FINDMYPWN**") was deployed against iOS 15 beginning in June 2022. It also appears to be a two-step exploit; the first step targets the iPhone's Find My feature, and the second step targets iMessage.*

- *We shared forensic artifacts with Apple in October 2022, and additional forensic artifacts regarding **PWNYOURHOME** in January 2023, leading Apple to release several security improvements to HomeKit in iOS 16.3.1*

- *Once we had identified **FINDMYPWN** and **PWNYOURHOME**, we discovered traces of NSO Group's first 2022 zero-click ("**LATENTIMAGE**") on a single target's phone. This exploit may also have involved the iPhone's Find My feature, but it utilizes a different exploit chain than **FINDMYPWN**.*

And here's the final comment that I wanted to share:

- *For a brief period, targets that had enabled iOS 16's Lockdown Mode feature received real-time warnings when **PWNYOURHOME** exploitation was attempted against their devices. Although NSO Group may have later devised a workaround for this real-time warning, we have not seen **PWNYOURHOME** successfully used against any devices on which Lockdown Mode is enabled.*

The first thing Apple notes when they're talking about the limitations imposed by Lockdown Mode is: *"Most message attachment types will be blocked, other than certain images, video and audio. Some features, such as links and link previews, will be unavailable."* We know from our examination last week of FORCED ENTRY that the way entry was forced was by sending the target a PDF with the .GIF file extension, which caused iMessage to attempt to render a very cleverly manipulated JBIG2 image that had been embedded in the PDF. It seems pretty certain that with Lockdown Mode, Apple has switched to a "Default Deny" with highly selective "Allows." So FORCED ENTRY would likely have also been nipped in the bud. The trouble with something like Lockdown Mode is that to be effective it really does need to be restrictive since, as we've seen, exploits are everywhere. That might annoy the people who need it most, enough for them to turn it off due to its interference with the things they need to do.

And we have fresh evidence that countries are busily using these patently illegal tools

NCSC stands for the National Cyber Security Centre in the UK, which is exactly what it sounds like. Last Wednesday, the Center published a report titled "Cyber experts warn of rising threat from irresponsible use of commercial hacking tools over the next five years." I've read the report and some of their conclusions serve as a useful reality check. Here are some selected pieces from the report:

The commercial proliferation of cyber tools and services lowers the barrier to entry to state and non-state actors in obtaining capability and intelligence that they would not otherwise be able to develop or acquire.

The sophistication of some commercial intrusion cyber products and services can almost certainly rival the equivalent capabilities of some state-linked Advanced Persistent Threat (APT) groups. The bulk of the commercial cyber sector is highly likely focused on satisfying domestic state demand from law enforcement and government agencies. However, over the last decade, a growing number of enterprises have emerged offering a range of products and services to global customers. They include off-the-shelf capability (Hacking-as-a-Service), bespoke hacking services (hackers-for-hire), and the sale of enabling capabilities such as zero-day exploits and tool frameworks.

Over the last ten years, at least 80 countries have purchased commercial cyber intrusion software, or spyware. For dozens of states without a skills base, the commercial sector is almost certainly transformational, allowing cost-effective access to capability that would otherwise take decades to develop.

While products vary in capability and application, commercially available spyware for mobile

devices can offer the ability to read messages, listen to audio calls, obtain photos, locate the device and remotely operate the camera and microphone. Some states are likely to procure multiple commercial cyber tools to meet their requirements. Devices can be compromised in a number of ways, including phishing, but also 'zero-click' attacks which do not require user interaction, making it more difficult for victims to mitigate.

While these tools have been used by states against law enforcement targets, spyware has almost certainly been used by some states in the targeting of journalists, human rights activists, political dissidents and opponents and foreign government officials. This is almost certainly happening at scale, with thousands of individuals targeted each year. While current products focus on mobile devices and intelligence gathering, as the sector grows and demand increases, products and services will likely diversify to meet demand.

Hacker-for-hire groups carry out cyber activity for paying clients. As well as providing information of traditional espionage value to states, hackers-for-hire are also reportedly used for legal disputes, intellectual property theft, insider trading, and the theft of other private data. Hackers-for-hire differ in skill and capability, ranging from low-level cyber crime activity to technically complex and effective network compromises that may go undetected. Some groups operate in criminal circles, some portray themselves as commercial companies, and others operate anonymously. Hacker-for-hire groups that focus on stealing information use phishing and other social engineering attacks, exploits against publicly reported vulnerabilities in computer networks, and sometimes zero-day attacks to compromise victims. The greatest threat comes from higher-end, hacker-for-hire groups, whose abilities and impact are similar to those of capable state actors. Hackers-for-hire pose a potential corporate espionage threat against organizations and individuals with privileged or valuable confidential information in multiple sectors.

While less-skilled and cyber criminal hackers-for-hire almost certainly carry out Denial of Service (DoS) attacks for a fee to temporarily disrupt a target website or server on a customer's behalf, additional law enforcement attention probably deters higher skilled hackers-for-hire from conducting destructive or disruptive operations. However, a growing market and the extra financial incentive raise the likelihood of hackers-for-hire accepting this type of tasking over the next five years. Hackers-for-hire also raise the likelihood of unpredictable targeting or unintentional escalation through attempts to compromise a wider range of targets, particularly those seeking valuable information to sell on, as opposed to 'working to order'. It is likely that potentially significant financial rewards incentivise state employees or contractors with cyber skills to become hackers-for-hire, risking the proliferation of cyber techniques from state to non-state actors.

Historically, underground criminal markets have facilitated the exploit trade. Since the early 2000s, a lucrative market for zero-day exploits has emerged in the commercial space. The large sums of money involved for critical zero-day exploits for commonly used systems and processes mean opportunities for profit are significant and have driven commercialisation.

Critical zero-day exploits and vulnerabilities are almost certainly transformational to actors with the skills to make use of them. States, or commercial cyber intrusion companies providing

products to states, are the dominant customers of the commercial zero-day market and are highly likely to remain so for the next five years. The growth of the commercial sector facilitating this trade has likely increased the number of states able to access critical zero-day capability, directly or indirectly.

Some well-funded cyber crime groups have highly likely purchased lower priced zero-day exploits for less well-used systems from underground exploit marketplaces. However, purchasing high-cost, critical zero-day capability from the commercial marketplace is unlikely to appeal to most cyber crime groups. Financial motivation makes it more likely that they prioritize lower-cost exploits developed from disclosed zero-day vulnerabilities, albeit as early as possible after disclosure to maximize the number of unpatched systems they can target.

Customisable tool frameworks are developed by cyber security software developers to emulate threat activity to enable penetration testing of networks. They are usually sold under license, but some are also publicly available or available in versions where the license has been removed. These frameworks are being used or repurposed by state and non-state actors; highly likely enabling a cost-effective uplift in cyber capability. It is highly likely that their constant evolution and the ability of actors to customize and repurpose these frameworks means widespread misuse of these frameworks will almost certainly continue over the next five years.

State and non-state actors also have access to capability developed and sold for cyber crime. In recent years, cyber crime marketplaces have grown and become increasingly professionalized, in part driven by demand from ransomware actors. One example is Malware-as-a-Service (MaaS), which is a service that provides use of malware, eliminating the need to create and develop the software as well as reducing the knowledge threshold required to operate the malware. Offering these services as a package is attractive to less skilled cyber criminals and as such has almost certainly expanded the number of victims.

And looking to the future, the report concluded:

Over the next five years:

- *Increased demand, coupled with a permissive operating environment, will almost certainly result in an expansion of the global commercial cyber intrusion sector, driving an increased threat to a wide range of sectors.*
- *It is almost certain there will be further high-profile exposures of victims against whom commercial cyber tools or hacker-for-hire operations have been deployed.*
- *Oversight of the commercial intrusion cyber sector will almost certainly lack international consensus, be difficult to enforce and subject to political and commercial influence.*
- *However, it is likely that many commercial cyber companies will be incentivised to vet and limit their customer bases, should effective oversight and international norms on development and sale of commercial cyber capability emerge.*

Last week we took a deep dive to appreciate the insane level of effort that the NSO Group's spyware developers exhibited in their successful effort to create another zero-click exploit of iPhones. Google's Project Zero researchers who reverse engineered this work reported that the sophistication of what they discovered terrified them. It was as if they had discovered alien technology lurking within a terrestrial device.

A growing black market for ChatGPT accounts

The best way for me to introduce this next topic is for me to just read what Check Point Research posted last week. Their headline was: *"New ChatGPT4.0 Concerns: A Market for Stolen Premium Accounts"* They wrote...

Since December 2022, Check Point Research (CPR) has raised concerns about ChatGPT's implications for cybersecurity. Now, Check Point also warns that there is an increase in the trade of stolen ChatGPT Premium accounts, which enable cyber criminals to get around OpenAI's geofencing restrictions to obtain unlimited access to ChatGPT.

The market of account takeovers (ATOs), stolen accounts to different online services, is one of the most flourishing markets in the hacking underground and in the dark web. Traditionally this market's focus was on stolen financial services accounts (banks, online payment systems, etc.), social media, online dating websites, emails, and more.

Since March 2023 [last month – so this is new behavior]*, Check Point sees an increase in discussion and trade of stolen ChatGPT accounts, with a focus on Premium accounts, including:*

- *Leak and free publication of credentials to ChatGPT accounts*
- *Trade of premium ChatGPT accounts that were stolen*
- *Bruteforcing and Checkers tools for ChatGPT – tools that allow cybercriminals to hack into ChatGPT accounts by running huge lists of email addresses and passwords, trying to guess the right combination to access existing accounts.*
- *ChatGPT Accounts as a Service – dedicated service that offers opening ChatGPT premium accounts, most likely using stolen payment cards.*

Why is the market of stolen ChatGPT account on rise and what are the main concerns?

As we wrote in previous blogs, ChatGPT imposes geofencing restrictions on accessing its platform from certain countries (including Russia, China and Iran). Recently we highlighted that utilizing the ChatGPT API allows cybercriminals to bypass different restrictions, as well as use of ChatGPT's premium account.

All this leads to an increasing demand for stolen ChatGPT accounts, especially paid premium accounts. In the dark web underground, where there is a demand – there are smart cybercriminals ready to take advantage of the business opportunity.

Meanwhile, during the last few ***weeks*** *there have been discussions of ChatGPT's privacy issues, with Italy banning ChatGPT and Germany considering banning it as well. We highlight*

another potential privacy risk of this platform. ChatGPT accounts store the recent queries of the account's owner. So when cybercriminals steal existing accounts, they gain access to the queries from the account's original owner. This can include personal information, details about corporate products and processes, and more.

During the last month, Check Point observed an increase in the chatter in underground forums related to leaking or selling compromised ChatGPT premium accounts.

Check Point then goes on to detail the specific workings of the underground dark web, account brute forcing, selling guaranteed access, etc. But I thought the interesting takeaway was that access to ChatGPT has become something of value and anytime that happens there will be dark forces at work to subvert its access.

Decommissioned Corporate Routers Leak Secrets

ESET made a very interesting observation in their posting last Tuesday, titled: *"Discarded, not destroyed: Old routers reveal corporate secrets"* Get a load of what they wrote:

Taking a defunct router out of an equipment rack and sliding in a shiny new replacement is probably an everyday occurrence in many business networking environments. However, ***the fate of the router being discarded*** *should be as important, if not more so, than the smooth transition and implementation of the new kit in the rack. Unfortunately, this appears often not to be the case.*

When the ESET research team purchased a few used routers to set up a test environment, there was shock among team members when they found that, in many cases, previously used configurations had not been wiped…and worse, the data on the devices could be used to identify the prior owners along with the details of their network configurations.

This led us to conduct a more extensive test, purchasing more used devices and adopting a simple methodology to see if data still existed on the devices. A total of 18 routers were acquired, one was dead on arrival, two were a mirrored pair so we counted them as a single unit; after these adjustments, we discovered configuration details and data on over 56% of the devices.

In the wrong hands, the data gleaned from the devices – including customer data, router-to-router authentication keys, application lists, and much more – is enough to launch a cyberattack. A bad actor could have gained the initial access required to start researching where the company's digital assets are located and what might be valuable. We are all likely aware what comes next in this scenario.

The change in recent years to the methods used by bad actors to conduct cyberattacks on businesses for the purposes of monetization is well documented. Switching to a more advanced persistent threat style of attack has seen cybercriminals establishing an entry point and a foothold into networks. They then spend time and resources conducting sophisticated extraction of data, exploring methods to circumvent security measures, and then ultimately

bringing a business to its knees by inflicting a damaging ransomware attack or other cyber-nastiness.

The initial unauthorized incursion into a company network has a value: the current average price for access credentials to corporate networks, according to research by KELA Cybercrime Prevention, is around ***$2,800****. This means that a used router purchased for a few hundred dollars, which without too much effort provides network access, could provide a cybercriminal with a significant return on investment. That's assuming they just strip the access data and sell it on a dark web market, as opposed to launching a cyberattack themselves.*

A concerning element of this research was the lack of engagement from companies when we attempted to alert them to the issue(s) of their data being accessible in the public domain.

Some were receptive to the contact, a few confirmed the devices had been passed to companies for secure destruction or wiping **[whoops!]** *– a process that had clearly not taken place – and others just ignored the repeated contact attempts.* **[wow.]**

The lessons that should be taken from this research are that any device leaving your company needs to have been cleansed, and that the process of cleansing needs to be certified and regularly audited to ensure your company's crown jewels are not being openly sold in public secondhand hardware markets.

We have published the details – well, all but the companies' names and data that would make them identifiable – in a white paper. The white paper also contains some guidance on the process that should be followed, including references to ***NIST special publication 800.88r1, Guidelines for Media Sanitization****. We strongly recommend reading the details and using our findings as a nudge to check the process in your own organization, to ensure no data is unintentionally disclosed.*

ESET's white paper is titled: *"How I (could've) stolen your corporate secrets for $100"*
https://www.welivesecurity.com/wp-content/uploads/2023/04/used_routers_corporate_secrets.pdf
SP 800-88 Rev. 1 — Guidelines for Media Sanitization:
https://csrc.nist.gov/publications/detail/sp/800-88/rev-1/final

In that white paper they provided a summary breakdown of what they found:

- *22% contained customer data*
- *33% exposed data allowing third-party connections to the network*
- *44% had credentials for connecting to other networks as a trusted party*
- *89% itemized connection details for specific applications*
- *89% contained router-to-router authentication keys*
- *100% contained one or more IPsec/VPN credentials, or hashed root passwords*
- *100% had sufficient data to reliably identify the former owner/operator*

I'll note that through 2023, so far, I've been having some very similar experiences of my own. I've been purchasing specific old drives from eBay when a SpinRite tester reports that they've

seen some weird behavior from an old drive. As SpinRite is running, one of its more popular screens flashes up snapshots of the data obtained from the drive it's scanning. The drives I've been purchasing second hand from eBay are **not** empty. NTFS file system metadata has a particular look, and these drives have been full of it. I have no interest in the contents of those drives beyond watching SpinRite recover whatever data they might contain. But they contain someone's data.

My experience, what ESET discovered in retired corporate routers, and even Check Point's previous observation about the value of stealing anything that can be stolen, should all serve as a potent reminder that, unfortunately, we all live in a world with plenty of bad guys who will take advantage of any opportunity to gain at others' expense.

Jaguar Tooth: Cisco router vulnerabilities

While we're on the topic of routers, let's take a look at last week's report from the UK's NCSC regarding "Jaguar Tooth" a Cisco router targeted malware. This serves as a perfect case study:

Jaguar Tooth is a system of backdoor Trojan malware deployed via exploitation of a long since patched SNMP vulnerability CVE-2017-6742. This vulnerability was first announced by Cisco on the 29th June 2017 when updated and repaired software was made available. Cisco's published advisory included details of workarounds, including through limiting access to SNMP from trusted hosts only (imagine that!) or by disabling the several vulnerable SNMP API branches, known as MIBs.

This amounts to another of those issues I so often have about policies vs mistakes.

SNMP — the "Simple Network Management Protocol" — is, essentially, a network API, a very powerful network API, which allows for the complete configuration state querying and configuration management of SNMP-capable networked devices. The point is, it should NEVER be publicly exposed to the wider Internet. It is meant to be used on the internal Intranet for internal management. And if by some weird network configuration need, a router's SNMP traffic does need to transit the public Internet, then it would certainly only ever need to be to or from a specific single targeted remote public IP. Never ALL public IPs — there's no conceivable reason for a router's SNMP service to be globally available. And yes, SNMP has an authentication layer. But it's old and its lame and it's barely adequate for the purpose of keeping insiders out — let alone outsiders. If external SNMP packets cannot reach the SNMP service, then vulnerabilities in that service will never become an issue.

So, again, policies versus mistakes. Mistakes happen by mistake. Okay. But policies happen by policy – in other words, on purpose. Mistakes don't need forgiveness, and policies don't deserve any. I don't mean to harp on this, but to me this delineation seems important and it is too often confused. Anyone who's responsible for any Cisco corporate router in 2023, which is still running a version of Cisco's OS from 2017 should be immediately summarily and disgracefully discharged from their responsibilities and their employment. It's unconscionable. But we also know that, unfortunately, many such routers will nevertheless exist.

So what do we know about this specific problem?

The vulnerability enables a stack-based buffer to be overflowed, enabling control of the instruction pointer which can be used to gain remote code execution. This exploit uses Return Oriented Programming (ROP) to overwrite operating system memory and incrementally deploy the malware code over hundreds of iterations.

The vulnerable function targeted by this exploit is reached using the SNMP Object Identifier (OID) which corresponds to alpsRemPeerConnLocalPort. By appending additional bytes to the end of the OID, a stack-based buffer can be overflowed.
One of the side-effects of this vulnerability is that any ASCII characters in the additional OID bytes are converted to uppercase, which constrains what data can be written and where.

Jaguar Tooth is deployed by writing custom shellcode to memory which can be used to write an arbitrary 4-byte value to any specified address. This shellcode is then called repeatedly to incrementally write Jaguar Tooth into memory. Once the Jaguar Tooth payloads have been copied into memory, they are individually executed by overflowing the return address of the vulnerable function with their location in memory.

Once Jaguar Tooth is running, it uses TFTP, the Trivial File Transfer Protocol, to exfiltrate pretty much everything the router knows about all of the peers that touch it. The router's ARP table is dumped to obtain the MAC addresses and IPs of all internal machines. And, of course, the bad guys now have a foothold in a border router, able to run whatever they choose and from there on it's all going to be pretty much bad news.

As I noted earlier, Cisco responsibly updated their IOS — Internetwork Operating System — posting the vulnerability disclosure back on June 29th, 2017, rating it High Severity with a CVSS of 8.8. But, of course, if a router is never going to be updated, no matter what, nothing that Cisco might do could possibly have any effect... short of arranging to never have made this implementation mistake in the first place. But no one in our industry appears to have figured out how to do that so far.

So, a malicious actor group known as APT28 has been detected actively conducting reconnaissance and deploying their malware on the world's routers which are still running that vulnerable version of IOS from 2017 when they also have SNMP exposed. Somehow, we need to figure out how to do better.

But until then, and probably always, we're really seeing the growing need for a position within an organization to be defined as responsible for nothing more than continuously inventorying and managing the patch level of all of an organization's equipment... no matter how obscure or dusty the closet may be. It may not be a glamorous job, but boy is it needed.

Security Research Legal Defense Fund
It’s time for some happy news. Something known as the *“Security Research Legal Defense Fund”* is in the process of being created and it is what its name suggests. The organization's website domain is also its name dot ORG with no spaces:
So: https://www.securityresearchlegaldefensefund.org/

They explain themselves in one long line: *"We aim to help fund legal representation for persons who face legal issues due to good faith security research and vulnerability disclosure in cases that would advance cybersecurity for the public interest."*

They break this down into three statements. First is their mission:

> *The Security Research Legal Defense Fund ("the Defense Fund") will be a nonprofit organization whose mission is to promote social welfare by providing financial assistance for legal representation of good faith security researchers and vulnerability disclosure.*

For Background, they say:

> *Society depends on secure digital communications and devices, but cyberattacks and system failures increasingly endanger physical safety, consumer privacy, and the operation of critical services.*
>
> *The public benefits when security vulnerabilities in software and systems are discovered and fixed before malicious actors can exploit them. In many instances, individuals have acted independently and in good faith to find and report vulnerabilities for mitigation, thereby strengthening the cybersecurity of products and services for the good of the community.*
>
> *While recognition from governments and businesses of the value of good faith security research and vulnerability disclosure is growing, individuals continue to meet with legal threats when their vulnerability research and disclosures are unwelcome or misunderstood. Such threats can ignore individuals' rights or misconstrue facts, creating a chilling effect on beneficial security research and vulnerability disclosure, especially for individuals without the resources to finance legal counsel.*

And, finally, under "How It Works" they explain:

> *The Security Research Legal Defense Fund may donate to good faith security researchers' choice of counsel to represent them in defending against claims related to good faith security research and vulnerability disclosure. The Defense Fund does not provide direct legal representation at this time. The organization's Board of Directors will consider potential grantees and vote on distribution of funds.*
>
> *To help ensure funds are used in the public interest, the recipients of legal defense funds would be required to meet eligibility criteria. The eligibility criteria is subject to revision by the Board, and aims to reflect alignment with legally accepted definitions of "good faith security research."*
>
> *The eligibility criteria to apply for grants from the Defense Fund is anticipated to include:*

- *The grantee demonstrates financial need;*
- *Funds donated from the Security Research Legal Defense Fund would go towards representation in legal matters related to good faith security research or vulnerability disclosure, and not such illegal behavior as extortion;*
- *The "good faith security research or vulnerability disclosure" was performed for the purpose of good faith testing, investigation, correction, or disclosure of a security flaw or vulnerability, was carried out in a manner designed to avoid harm to individuals or the public, and the information derived from the activity was intended to be used primarily to promote the security or safety of computers or software, or those who use such computers or software; and*
- *Board approval.*

Through the years, here, we've talked about this problem is well meaning, typically amateur hackers not backed by an organization, attempting to inform an organization of some significant problem they've stumbled upon and identified, only to have the organization's management freak out and aim law enforcement and their attorneys at the hapless hacker. So this seems like a terrific backstop for such situations.

I was curious to learn more about where this came from, so I did a bit of digging around and found that SC Magazine had what I was looking for. I've edited a bit of what they wrote:

Google and other companies will develop and stand up a pair of new initiatives that will provide policy guidance to governments and legal protection to security researchers engaged in "good faith" vulnerability research and disclosure, while the tech giant also said it would formalize an internal policy to be publicly transparent when bugs in Google products are exploited in the wild.

The moves include the establishment of an industry-led Hacking Policy Council, which would be designed to bring "like minded organizations and leaders who will engage in focused advocacy, new policies and regulations, support best practices for vulnerability management and disclosure, and do not undermine our user's security," as well as a planned nonprofit that would fund legal costs for security researchers who are sued or prosecuted while conducting vulnerability research and disclosure.

The council will include representatives from bug bounty firms HackerOne, BugCrowd, Intigriti and Luta Security, as well as Intel and Venable, a law firm that specializes in cybersecurity law and policy matters.

Charley Snyder, head of security policy at Google, when asked how the council chose its initial membership, said: "I think it's very much a coalition of the willing. There was no real criteria for membership. This is a fairly specialized area of policy, and these companies are ones that are really invested in getting it right."

The formation of the council comes at a time when the United States and other nations are showing an increased willingness to regulate the cybersecurity choices of businesses and other entities to prevent cyberattacks from significantly disrupting or spreading through a particular sector, critical infrastructure and other essential services.

Or, in other words, as we've recently been noting, the cyber security terrain is becoming increasingly litigious.

> *The use of existing or future regulatory authority was a key pillar of the Biden administration's national cybersecurity strategy, and agencies like the Securities and Exchange Commission, the Transportation Security Administration and the Environmental Protection Agency have since come out with a raft of sector-specific cybersecurity regulations or proposals over the last month.*
>
> *The other announced initiative is a legal defense fund for security researchers who are sued or prosecuted for pursuing "good faith research in cases that would advance cybersecurity for the public interest." Representatives from Google told SC Media they will provide seed money for the fund but it will be managed as a separate, non-profit 501c3 entity.*
>
> *At an event Thursday, Harley Geiger, a cybersecurity attorney at Venable, said incidents "such as when the governor of Missouri threatened a reporter for telling a state agency about a vulnerability in his website" would be among the cases the fund will be set up to support. Geiger said the fund does not plan to provide direct legal representation or services to affected researchers "at this time."*
>
> *The Department of Justice under President Biden has made it official policy to avoid prosecutions of "good faith" security research under the Computer Fraud and Abuse Act. But digital civil liberties organizations continue to have questions about how law enforcement agencies will define such efforts, and apart from that, researchers and journalists have also faced the threat of lawsuits from private sector companies and organizations when they disclose or report on vulnerabilities in their products.*
>
> *Katie Moussouris, the CEO and founder of Luta Security said: "Right now, we have a lot of regulations that were written in a different era when there was not a nuanced understanding of the hand-in-hand relationship between vulnerability discovery and malicious hacking prevention. I think that it's very difficult to get a lot of these regulations to be unwound to a productive place [and] I think there is a lot of room for improvement [and] disambiguation of intent."*

A quick Firefox fix: Last week Firefox users moved to v112.0.1 to fix exactly one problem. Mozilla wrote: "Fixed a bug where cookie dates appear to be set in the far future after updating Firefox. This may have caused cookies to be unintentionally purged." I didn't notice any problem and I am once again, after using Google and Bing for a spell, happily back to the Firefox fold, choosing to use it as my primary browser.

Kubernetes security audit: NCC Group has concluded and published a new security audit of the Kubernetes automation platform. Nothing significant was found.

Google Chrome zero-day: Google has released Chrome v112.0.5615.137/138 that fixes eight security flaws, including a new zero-day (CVE-2023-2136). The zero-day was discovered by Google's TAG team and came after the team patched another zero-day last week. That one was abused by a surveillance vendor.
https://chromereleases.googleblog.com/2023/04/stable-channel-update-for-desktop_18.html

An End-to-End Encryption Proposal

We have, of course, been covering the fascinating and escalating debate over the presence of ubiquitous end-to-end encryption which took another step with the UK's Online Safety Bill which is currently winding its way through the United Kingdom's legal system, but is on its way to becoming law in the UK. As we know, this is the legislation that's being promoted as a means of protecting children from online threats of all sorts by requiring secure messaging providers to somehow arrange to monitor and filter the images, videos, audio and textual communications of their entire user base, regardless of whether individual users are suspected of illegal behavior.

While assembling today's podcast, I encountered two opposing open letters which I'll share here in a moment. What's surprising is that after reading and placing these open letters into the podcast's show notes, I was summarizing the current situation and working through the dilemma... and I may have actually come up with a workable solution to this whole encryption mess. I'm not kidding. It's kind of perfect. Okay, but first things first:

Last week's news is that the CEO's of the secure messaging firms have collectively authored and co-signed an open letter to the UK Government. Represented, were the heads of Element, Session, Signal, Threema, Viber, WhatsApp and Wire. Since this open letter contains a few juicy bits, I want to share what the heads of today's secure messaging companies wrote to the UK:

To anyone who cares about safety and privacy on the internet. **[Just a bit loaded]**

As end-to-end-encrypted communication services, we urge the UK Government to address the risks that the Online Safety Bill poses to everyone's privacy and safety. It is not too late to ensure that the Bill aligns with the Government's stated intention to protect end-to-end encryption and respect the human right to privacy.

Around the world, businesses, individuals and governments face persistent threats from online fraud, scams and data theft. Malicious actors and hostile states routinely challenge the security of our critical infrastructure. End-to-end encryption is one of the strongest possible defenses against these threats, and as vital institutions become ever more dependent on internet technologies to conduct core operations, the stakes have never been higher.

As currently drafted, the Bill could break end-to-end encryption, opening the door to routine, general and indiscriminate surveillance of personal messages of friends, family members, employees, executives, journalists, human rights activists and even politicians themselves, which would fundamentally undermine everyone's ability to communicate securely.

The Bill provides no explicit protection for encryption, and if implemented as written, could empower OFCOM [the UK's communications regulator] *to try to force the proactive scanning of private messages on end-to-end encrypted communication services - nullifying the purpose of end-to-end encryption as a result and compromising the privacy of all users.*

In short, the Bill poses an unprecedented threat to the privacy, safety and security of every UK citizen and the people with whom they communicate around the world, ***while emboldening hostile governments who may seek to draft copy-cat laws.***

Proponents say that they appreciate the importance of encryption and privacy while also claiming that it's possible to surveil everyone's messages without undermining end-to-end encryption. ***The truth is that this is not possible.***

We aren't the only ones who share concerns about the UK Bill. The United Nations has warned that the UK Government's efforts to impose backdoor requirements constitute ***"a paradigm shift that raises a host of serious problems with potentially dire consequences"****.*

Even the UK Government itself has acknowledged the privacy risks that the text of the Bill poses, but has said its ***"intention"*** *isn't for the Bill to be interpreted this way.* [WHAT?!?! How are we supposed to ***"interpret"*** it?!?! We're not interpreting it, we're reading it!]

[And here's the money shot...]

Global providers of end-to-end encrypted products and services cannot weaken the security of their products and services to suit individual governments. There cannot be a "British internet," or a version of end-to-end encryption that is specific to the UK.

The UK Government must urgently rethink the Bill, revising it to encourage companies to offer more privacy and security to its residents, not less. Weakening encryption, undermining privacy, and introducing the mass surveillance of people's private communications is not the way forward.

Signed by those who care about keeping our conversations secure.

Okay. So there's the open letter from the encryption **providers** who argue, convincingly, I think, that forcing surveillance capability into all communications is not a workable idea.

An open letter from the Virtual Global Taskforce

So next we have a second open letter published last Wednesday by a group known as the Virtual Global Taskforce. They describe themselves as an international alliance of 15 law enforcement agencies. I was a bit suspicious because the chair of the organization is the UK's National Crime Agency. So I was wondering how global they were. But Wikipedia knows all about them and explains:

The Virtual Global Taskforce (VGT) is a group of law enforcement agencies from around the world who operate together to stop online child sex abuse. The VGT is made up of the following organizations:

- *Australian Hi-Tech Crime Centre / Australian Federal Police (AFP)*
- *Child Exploitation and Online Protection Centre (United Kingdom)*
- *Colombian National Police*
- *Cybercrime Coordination Unit Switzerland (CYCO)*
- *Dutch National Police*
- *Europol*
- *Interpol*

- *Italian Postal and Communication Police Service*
- *Korean National Police Agency*
- *Royal Canadian Mounted Police*
- *New Zealand Police*
- *Ministry of Interior for the United Arab Emirates*
- *Philippine National Police*
- *U.S. Immigration and Customs Enforcement (ICE) of the US's DHS.*

Okay. So this group has collectively authored and sent an open letter to Meta asking the company to reconsider adding end-to-end encryption features to Facebook and Instagram. The letter argues, of course, that this would hinder their own and Meta's efforts to fight the proliferation of CSAM on the platform. Here's what they said:

The Virtual Global taskforce is calling for all industry partners to fully appreciate the impact of implementing system design decisions that result in blindfolding themselves to child sexual abuse (CSA) occurring on their platforms, or reduces their capacity to identify CSA and keep children safe.

It is time to confront these concerns and make tangible steps towards possible solutions that we know exist.

The Virtual Global Taskforce is an international alliance of 15 dedicated law enforcement agencies, of which the National Crime Agency is the chair, working alongside Affiliate members from private industry and non-governmental organisations to tackle the threat of child sexual abuse (CSA).

The VGT issued its first position statement on end-to-end encryption (E2EE) in 2021. This statement highlighted the devastating impact E2EE can have on law enforcement's ability to identify, pursue and prosecute offenders, when implemented in a way that affects the detection of CSA on industry platforms. It is important to update the VGT position on E2EE in the context of impending design choices by industry.

As outlined in our previous statement, there is no doubt that encryption plays an important role in safeguarding privacy, however this must be balanced with the importance of safeguarding children online.

The VGT encourages industry to respond and consider the following:

- *Only to implement platform design choices, including E2EE, at scale alongside robust safety systems that maintain or increase child safety.*

- *Where the child user base and risk is high, a proportionate investment and implementation of technically feasible safety solutions is paramount.*

The abuse will not stop just because companies decide to stop looking. We all have a role to play in protecting children in online spaces and we strongly urge industry partners to take

active steps toward this goal.

The scale of online CSA is increasing worldwide. The WeProtect Global Alliance have identified it as one of the most urgent and defining issues of our generation. The number of reports of CSA from industry continue to be staggering, but demonstrates the key role that industry plays both in protecting children online and in reporting cases to law enforcement for action.

The National Center for Missing and Exploited Children (NCMEC) received 29.3 million reports of suspected CSA in 2021, a 35% increase from 2020. Of this 29.3 million, over 29.1 million reports came from electronic service providers.

Although these reports result in a range of different outcomes globally, what is consistent is that they significantly contribute to positive outcomes for child safety. These figures demonstrate the current success of industry partners in detecting and reporting CSA occurring on their platforms, resulting in victims being identified and safeguarded.

Design and investment choices implemented in a way that interferes with the effectiveness of such safety systems threaten to undermine these successes which have been consistently built upon over previous decades.

The announced implementation of E2EE on META platforms Instagram and Facebook is an example of a purposeful design choice that degrades safety systems and weakens the ability to keep child users safe.

META is currently the leading reporter of detected child sexual abuse to NCMEC. The VGT has not yet seen any indication from META that any new safety systems implemented post-E2EE will effectively match or improve their current detection methods.

So, everybody's gearing up and staking out their positions. It's unclear what's going to happen. Legislation is probably going to be passed, since it's easy for politicians to write laws which tell others what they can and cannot do. But it's difficult to see any of the providers of end-to-end encryption backing down; especially not those like Telegram, Signal and Threema whose entire purpose is secure end-to-end encryption. Apple proposed a solution that would be minimally invasive, but the public freaked out over the idea of anything like a library of known child pornography being resident on their phones, and that sentiment is understandable. And Apple's solution would not handle the whole text messaging "grooming" problem.

So this all led me to revisit the question we touched upon once previously, which was whether some form of good old fashioned parental control might be the only answer. Perhaps we would need to decide that these social media devices are just too dangerous for children to have? And that led me to an interesting idea that I haven't seen suggested anywhere before:

Why don't we arrange to only monitor children?

The pending laws and legislation is changed to **only** require exactly what those governments claim is their reason and motivation for needing to compromise full end-to-end encryption...

Which is sexual abuse material (CSAM) content and behavior screening **for minors**. Then we implement that legislation with technology so that the devices children use in countries that require it are aware of the date when they will no longer be subject to monitoring for their own protection.

When any device is first set up and configured with an account, the setup process determines whether the user of this device resides within a country whose government has mandated the surveillance of minors. If not, that's the end of it. But if so, the setup process is then informed of whether the user is already an adult. If so, again, that's the end of it. But if the user is currently a minor in their local society, governed by laws which mandate the protection of minors online, the setup process asks the user's date of birth and the age at which they will no longer be considered a minor in their world. This sequence of steps sets and stores an immutable date which subsequently governs the behavior of all encrypted services available for the device. Encrypted services query for a binary value: whether or not its user requires the protection provided by side-channel content moderation.

While users are young, any government mandated surveillance will be conducted in the background without interfering with the use of any applications. It will be entirely transparent to its young users. But on the day of their birthday, when they reach the age of majority, all such background side-channel surveillance automatically terminates – in full compliance with the laws governing their use of encrypted services in their society.

And, importantly, this solution means that no user who is already an adult – none of us, for example – will ever be subjected to this monitoring.

Think about all the problems this solves:

Children don't lose any functionality. Everything works for them as it always has. Yes, sure, in the margins they're sacrificing some of their privacy in the interest of their protection from online predation. But only while it's in their best interests to be protected. As soon as it's no longer needed, it disappears. And since there's no observable effect from its presence, there's no great pressure to cheat the system. Children are never inconvenienced. Everything works perfectly for them and the side-channel monitoring is completely invisible.

Parents can take some relief in knowing that, whatever it is their kids are doing online, it's being monitored for their safety while preserving as much of their privacy as possible. So parents, who are in the position to oversee the setup of this system – in compliance with their local laws – are able to enforce its presence.

Adults, who are not endangered by online exploitation, enjoy the privilege of truly private unmonitored end-to-end encryption without any fears of big brother eavesdropping. The fact that adults are **never** monitored dispels the worries about eventual government overreach and the presence of hidden government surveillance agendas. Only children are monitored.

The online slimeballs who seek to take advantage of youthful trust and innocence **know** that all of their communications with any underage targets **are being monitored**. So that hopefully pours some cold water where it may do some good.

The concern of whether such surveillance might be a slippery slope, and whether governments are actually using "but think of the children" as a stalking horse to mask their real interest in perpetrating more widespread surveillance is resolved by this. **No adult is monitored** – only young users whose electronic devices are aware of their monitoring cutoff date are protected.

If governments have secret intentions to expand this monitoring beyond sexual exploitation of minors, then that's fine too... but whatever they do, it will only work on kids.

The monitoring cutoff date system could be entirely local to the device as I described above, setup under parental supervision when the device is first brought online and never subject to change. Or it might be set by the device's service provider, such as a cell provider, or by the device's account provider, such as Apple, Google, Samsung, etc. In the future, if governments require some form of oversight verification that minors are being protected, that, too, could be implemented.

But the crux of this idea is clear. We've come to loggerheads and are approaching an impasse because both sides have been taking absolute all-or-nothing positions. Using technology, a compromise is possible that should satisfy everyone. Governments and law enforcement agencies say that they want to monitor children for their own protection. Fine, that can be arranged. Adults are adamant that they do not want to ever be monitored. Fine. They won't ever be. Everyone worries that governments have a hidden agenda for this monitoring, this makes that impossible.

I've been thinking about this since it first occurred to me yesterday and I can't find much to fault it.

The need to embed the date when surveillance will no longer be needed is new. But so what? New features are being added to our phones continually. If necessary during a transition period, or when a device does not yet offer the "Monitor Me" flag, the age determination could be individually distributed among encrypted service providers when accounts are created. But it would be cleaner to have this built into the device and queryable by encrypted apps.

If a device is shared by multiple children, the age of the youngest user would be chosen so that the youngest user remains protected and all remain in compliance with local laws.

In all of the coverage we've given of this mounting encryption standoff I've never seen any mention of something like this that appears to be a workable compromise. Both sides sound and appear to be absolute in their positions. But this would appear to offer a compelling middleground that would not be objectionable to either adults or pre-adults. And it feels like a compromise that even the encryption absolutists could live with – and they may have to if they with their services to remain legal where monitoring of minors is required by law.