An End-to-End Encryption Proposal

Description: 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 "Pegasus-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!

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-920.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-920-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about a couple of zero-click exploits discovered. NSO Group was selling them. Some of them might even get around lockdown mode on the iPhone. We'll also talk about why you should wipe your routers before you sell them or give them away. And speaking of routers, a flaw in Cisco's IOS that hasn't been patched in some older routers and why you should fix it now. It's all coming up next, a lot more, too, on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 920, recorded Tuesday, April 25th, 2023: An End-to-End Encryption Proposal.

It's time for Security Now!, yay, the show we wait all week for. In my case, all month for.

Steve Gibson: Guess who's back?

Leo: I'm here, kids.

Steve: Yes, indeed.

Leo: And I missed you, Uncle Steve.

Steve: Well, Leo, Ant and Jason did a great job of holding the fort down.

Leo: Thank you, guys. I really appreciate that.

Steve: The one piece of news that really had me thinking of you was the news that we covered in February of that ridiculous psychotherapy clinic in, I don't remember where, in Europe somewhere - Finland, I think - where a hacker got in and stole the records of 30,000 of their previous clients.

Leo: Oh, yeah, I remember.

Steve: And when the clinic didn't succumb to the ransomware threats, the hacker emailed all the clients, saying, "I've got all your data. I need 300 bucks from each of you."

Leo: Oh, god.

Steve: "Or I'm going to expose you to the Internet." Well, the good news is the CEO is convicted and in jail.

Leo: Oh, good.

Steve: So, yeah, I was thinking of you when we were...

Leo: They take that stuff seriously there.

Steve: Yeah, yeah.

Leo: I've got to tell you one thing, though, that's extremely annoying. You think these cookie banners are bad in the states?

Steve: Uh-oh.

Leo: Every page you go to it covers the page. You can't do anything until you see this big - it's much bigger, much more elaborate, lots more clauses. And every single place you go to you've got to click that. So they take that law seriously, too. And it's so stupid because what are you going to click but accept, accept, accept; right?

Steve: Yeah.

Leo: It's just dumb. Anyway, I'm glad to be back, and I'm so glad to see you. I did, now don't take this amiss, but I did, you know, we have an AI Leo now in our...

Steve: I saw that you asked ChatGPT how many - when would I be doing the last episode of Security Now!.

Leo: I'm just trying to do some planning.

Steve: It was quite wonderful.

Leo: Yeah. I said, because as you said, you've said this again and again, 999 is the last episode because you don't have four digits in you. And we're doing 920 today. So I asked, first I asked ChatGPT, in fact GPT-4, the official Open AI one, to tell me, well, when will Steve's last show be? It went on and on and on, "To calculate the date that's 79 weeks from April 25th, 2023, we follow these steps," and came up with the wrong freaking answer.

Steve: Well, and I love it, too, because then it begins moving month to month, trying to figure out how many days each month has. I mean...

Leo: And then it gives up.

Steve: Yeah, it's sort of amazing that it got, like, it thought it knew what it was doing.

Leo: It's bizarre.

Steve: And it is the point that is worth making. "60 Minutes" did a piece on all of this, I think it was Sunday before last. I mean, I think it was Scott Pelley who had - it was the first half of the show. And he was astonished by some of the things that this thing came up with, like they gave it, like, Hemingway's shortest writing, which was three two-sentence words, and asked it to expound on it or create a short story from it. And it did this amazing job.

Leo: Yeah. Well, that's what it does.

Steve: And then someone said, "And we'd like it as prose." And it turned it into a poem.

Leo: Sure.

Steve: Lorrie had tears running down her face because it was just so eloquent and amazing. And then, but the point is, then they said, "Give us your forecasts for the economy."

Leo: Oh.

Steve: And again, it produced this beautiful...

Leo: Very credible, yes.

Steve: ...fantastic-looking result. And then said, "For additional information, you might check these five reference texts," none of which were actual. It made them up.

Leo: It makes stuff up, yeah, because it's just fancy autocorrect. I did go to Wolfram Alpha, which can do math. And just in case anybody wants to know, unless we miss an episode because of a Best Of or whatever, October 29th, 2024, will be 999.

Steve: Well, and you actually keep the counter, you keep the counter going through those.

Leo: Yeah, I think we should - it should be accurate, yeah.

Steve: Well, yeah. And because, I mean, I've had my count off when I forgot to add one across the holidays.

Leo: Yeah, we fix it, yeah.

Steve: So it's like, uh-oh.

Leo: We try to get it, you know, so you have 52 episodes a year. And if that happens, it'll be, I'm sad to say, sorry ChatGPT, not August 30th, but, thank goodness, October 29th. It'll be Halloween weekend.

Steve: By that point you could probably just plug it in in place of me, Leo, and just let ChatGPT, you know...

Leo: Well, it kind of knows now. It's got a deadline. Yeah, get going. So what's coming up today?

Steve: Okay. So, oh, I've got something. I've got something, I think. So this is Episode 920 for - this is our last episode of April, believe it or not, Leo. You missed, like, most of...

Leo: I missed April, yeah.

Steve: Now, here we are, it'll be May before we know it. Not May Day, we're going to miss that by one, but that's good. Today's topic is An End-to-End Encryption Proposal. I think I may have solved this whole problem with government spying and all that, and the tension between problems.

Leo: Oh, that would be good.

Steve: That would be, yeah. Anyway, we're going to get to that. We're going to look at the past week's most interesting security news answers for the questions which have arisen, like whether Apple's Lockdown Mode actually does anything useful. How big is the market for commercial Pegasus-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? Whoops. And speaking of routers, what was the mistake that users of old Cisco routers really wish Cisco had not made, and whose fault is its exploitation today? What's the story behind the newly established Security Research Legal Defense Fund? And 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 anyone's going to be getting bored this week.

Leo: Yeah, I can't wait.

Steve: And of course we do have a great picture.

Leo: Yes, we do. Yes, we do. Picture time. This is hysterical, by the way.

Steve: Oh, okay. So yes. We have a picture. Thank you, one of our listeners, for tweeting this to me. Very valuable. This is a 16-port D-Link, barely recognizable router, or rather an unmanaged switch, yes.

Leo: It's a switch; right? A switch, yeah.

Steve: An unmanaged switch. If the guy in charge of the set of the Addams Family had been overpaid, I think this is...

Leo: This is a ceiling or the basement somewhere and just...

Steve: Oh, god.

Leo: Look at this. Ugh.

Steve: Yeah. So as I said, it's barely recognizable. I had to, like, squint to, okay, so we've got, you know, we have, on the upper side we have four - we have two groups of four ports, and I'm sure it's the same thing below. It looks like there is a little 4x4 grid of LEDs. And they're lit up. A lot of them are green there. So I gave this the caption, "By some miracle, it's still working. Don't touch it." Because, boy, I mean, yeah. So I don't

know where, probably like some office building, like sort of a low-end office facility where...

Leo: It's so dusty and cobwebby, though. I mean, that's disgusting.

Steve: Yeah. And why would one ever put this in the attic where it's, you know, I mean, you can just see it's like there's some sort of a heating duct passing by in the very far lower left corner. You sort of see the ribs of the duct. And many a spider has hopefully set up its operations here and thought, okay, I'm going to get lucky. It doesn't look like anything happened there except the web collected a lot of dust, you know, thus the Addams Family.

Leo: It's ridiculous.

Steve: So anyway, yes.

Leo: Oh, my.

Steve: Another great piece of technology that somehow is keeping the whole world online. If the Internet goes through there, then yeah, you're in danger.

Leo: Wouldn't that be funny if you unplugged it, and the whole Internet went down?

Steve: Yeah, exactly. Turns out that's the hub.

Leo: That's it.

Steve: That's the main crux, yeah. So last Tuesday, the forensic security research group Citizen Lab reported on three iOS 15 and 16 exploits attributed to Israel's NSO Group's Pegasus smartphone spyware system. And now, just as a reminder, Citizen Lab is at the University of Toronto's Munk School of Global Affairs and Public Policy. They've been doing some serious forensic work. Last week's main topic was ForcedEntry, which for our listeners and Leo, I'm bringing you up to speed, it was a fascinating example 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, who found a live sample of it and sent it to Project Zero to take a look at. They found it, Citizen Lab found it on a phone of a Saudi political activist.

So my eye was caught by Citizen Lab's mention just last week of the apparent successes of Apple's Lockdown Mode, which of course we previously described when Apple announced it. I think it was brand new in iOS 16. I think that's where it first appeared. As we know, it 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 their iOS, of Apple's iOS devices.

Last Tuesday they published an extensive description of several zero-click attacks they discovered, "they" being Citizen Lab, discovered being deployed against users of iOS 15

and 16. So not old, out-of-service version 13, or 12 and 13, 14. These are today's phones. I'm not going to get into all of the details of those, but their intersection with Lockdown Mode I thought was interesting. So to give you some sense, before we get to the intersection, they summarized their finds, the things they found, in seven bullet points which are short, and so I think worth sharing.

They said: "In 2022," so just last year, "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 16 zero-click exploit chains against civil society targets around the world." And just so everyone is on the same page here, "zero-click exploit" means the user does nothing. We covered one of them in detail last week where just the receipt of an iMessage, without any acknowledgment of any kind, was enough to take over the target's phone.

Leo: And I'm sure what this is covering was that Apple was patching these.

Steve: Yes, right.

Leo: I mean, very aggressively. I think they had three patches in a few weeks.

Steve: Right, right. And they've been like, again, they're doing everything they can. But it just seems that more problems keep being found.

Leo: It's aggressive, yeah.

Steve: Yes. So Citizen Lab said, speaking of these: "NSO Group's third and final known 2022 iOS zero-click, which we call 'PwnYourHome,' was deployed against iOS 15 and 16 starting in October of 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" - thus PwnYourHome - "and the second step targets iMessage." So somehow the exploit against HomeKit primes iMessage to then be exploitable by the second step in the two-step zero-click.

Leo: And that's not unusual. I mean, Pwn2Own we see multiple steps sometimes.

Steve: Chains.

Leo: Yeah, chains, yeah.

Steve: Chains, exactly, yeah. So then they said: "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" - thus FindMyPwn - "and the second step targets iMessage." They said: "We shared forensic artifacts with Apple in October of 2022, and additional forensic artifacts regarding PwnYourHome in January of this year, 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 iPhone. This exploit may also have involved the iPhone's Find My feature, but it utilizes a different exploit chain than FindMyPwn."

Okay. And so here's the final comment that I wanted, with all of that background, 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."

Leo: That's great.

Steve: Yes.

Leo: That's amazing.

Steve: Yeah. They said: "Although NSO Group may have later devised a workaround for this real-time warning, even so, we have not seen PwnYourHome successfully used against any devices on which Lockdown Mode is enabled." So props to Apple.

The first thing Apple notes when they're talking about the limitations that they impose when Lockdown Mode is enabled, they say: "Most message attachment types will be blocked, other than certain images, video, and audio." They said: "Some features, such as links and link previews, will be unavailable." And so of course we know from, as I said, our examination last week of ForcedEntry 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 then highly selective "Allows" being permitted. So ForcedEntry would likely have also been nipped in the bud.

The trouble with something like Lockdown Mode, unfortunately, is that to be effective it really does need to be a restrictive service, I mean, it needs to be restrictive. As we've seen, exploits are everywhere. And that might annoy the people who need it the most, enough for them to turn it off due to its interference with the things they need to do. So just this is really, of everything we've seen, Lockdown Mode hasn't been around for that long. But the ability to find the forensic evidence of attacks which are this targeted, as highly targeted as Pegasus is, it's a very rarified environment there. So this is neat feedback to have. And we have fresh evidence that countries are busily using these patently illegal tools, governments. Yeah, I know.

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." Which of course begs the question, what would "responsible use" be? It's like, okay, I mean, this is all illegal; right? And the countries, most of the countries, probably not all of them because they're certainly used in non-Democratic repressive regimes, but you have to imagine that a lot of the countries who said, oh, no, no, no, that's bad, it's like, okay, how much does that cost? They just don't seem to be able to say no.

Okay. Here are some selected pieces from the report, which I think provide some conclusions which serve as a useful reality check. So the UK's NCSC said: "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." I mean, and that really is the key; right? It's not like you need to have your own NSA-level capability anymore. You just go to the NSO Group and say, "How much?" and you get what you need.

They said: "The sophistication of some commercial intrusion cyber products and services can almost certainly rival the equivalent capabilities of some state-linked Advanced Persistent Threat groups. The bulk of the commercial cyber sector is highly likely focused on satisfying domestic state demand from law enforcement and government agencies." Right? I mean, the providers don't exist in a vacuum. They're fulfilling a need. They're selling this to people because people want it.

"However," they wrote, "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" - known as 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 past 10 years, at least 80" - eight zero - "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." Wow. "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." I mean, we're talking about a whole ecosystem which is emerging, on the DL, being sold to governments.

They write: "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 cybercrime 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 cybercriminal 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." In other words, the really high-end guys, they're not mucking around down in the less skilled areas of DDoS attacks. And those tend to be much more easily attributable, as we see.

"However," they said, "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." So that might be changing. "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.'" So they might be like out doing their own reconnaissance work, hacking companies just, you might say, on spec, as opposed to under contract. They said: "It's likely that potentially significant financial rewards incentivize state employees or contractors with cyber skills to become hackers-for-hire, risking the proliferation of cyber techniques from state to non-state actors." So as you might expect, these skills spread over time. They're not, you know, containment is not being maintained from where they originated.

"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 commercialization." And of course we've often talked about the likes of Zerodium, who are purchasing these exploits for resale. And there's no accountability. There's no sense for to whom they are selling them. But Zerodium makes no bones about it. It's like, yeah, we want to buy your exploits. We'll pay you dollars.

They said: "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 for 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 cybercrime 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 cybercrime 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." Of course this all follows from the things we're talking about on this podcast constantly.

"Customizable tool frameworks are developed by cybersecurity software developers to emulate threat activity to enable penetration testing of networks." And we were just talking about that with red team attacks. They said: "They're 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's 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 cybercrime. In recent years, cybercrime marketplaces have grown and become increasingly professional, 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 cybercriminals, and as such has almost certainly expanded the number of victims."

So, anyway, I thought they did a great job of sort of encapsulating everything that's going on there and the trends that all of the evidence points to as where we'll be headed in the future. And so they concluded with just four points, saying: "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. Second, 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.

"Third, oversight of the commercial intrusion cyber sector will almost certainly lack international consensus, be difficult to enforce, and subject to political and commercial influence." Right? I mean, the people who you would like to be providing the oversight are the customers. And they finished: "However, it is likely that many commercial cyber companies will be incentivized to vet and limit their customer bases, should effective oversight and international norms on development and sale of commercial cyber capability emerge." So, you know.

Last week we took a deep dive, as we know, into ForcedEntry, which led us to appreciate the insane level of effort that the NSO Group's spyware developers, whomever they are, and what we just shared suggests maybe they're not even NSO Group developers. They might have developed a zero-click independently and said, "How much is this worth to you?"

Leo: I thought the NSO Group disbanded. That's why I'm surprised to see their name again.

Steve: They did get pounded down, well, because their stuff was highly publicized.

Leo: Right.

Steve: But they're still around.

Leo: Can't kill cockroaches.

Steve: Exactly.

Leo: They come back again.

Steve: So anyway, what Google's Project Zero researchers who reverse-engineered that work reported was that the sophistication they discovered in that exploit terrified them. They used the word "terrified." It was as if they had discovered alien technology lurking within a terrestrial device. Which, you know, gives you some pause to note the obvious commercial value of this sort of technology. And now we're seeing, almost not surprisingly, a growing black market for ChatGPT accounts, which I thought, what?

The best way for me to introduce this next topic is to just read what Check Point Research posted last week. Their headline was: "New ChatGPT 4.0 Concerns: A Market for Stolen Premium Accounts." And they said: "Since December of 2022, Check Point Research has raised concerns about ChatGPT's implications for cybersecurity." Remember we talked about some of this before, where they used it both to try to reverse engineer for them some code, which it did not do a good job at; or writing code, which unfortunately we know it does do a good job at.

So they said: "Now, Check Point also warns that there is an increase in the trade of stolen ChatGPT Premium accounts, which enable cybercriminals to get around OpenAI's geofencing restrictions to obtain unlimited access to ChatGPT. The market for account takeovers" - generically 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, and so forth - social media, online dating websites, emails, and more."

They said: "Since March of 2023" - which was four weeks ago - "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, trading of Premium ChatGPT accounts that were stolen, brute forcing and Checker tools for ChatGPT" - meaning tools that allow for brute forcing - "that allow cybercriminals to hack into ChatGPT accounts, and ChatGPT Accounts as a Service, dedicated service that offers opening ChatGPT Premium accounts, most likely using stolen payment cards."

So why is the market of stolen ChatGPT account on the rise? What are the main concerns? They said: "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 using ChatGPT AI allows cybercriminals to bypass different restrictions, as well as use of ChatGPT's Premium account. All this leads to an increased demand for stolen ChatGPT accounts, especially paid Premium accounts. In the dark web underground, where there is a demand, there are smart cybercriminals already taking advantage of this business opportunity.

"Meanwhile, during the last few weeks there have been discussions of ChatGPT's privacy issues, with Italy banning ChatGPT, and Germany now considering a ban, 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." And Leo, on the first of the three podcasts where Ant co-hosted, we covered the story of Samsung's employees on three separate instances uploading, wanting to get the advantage of ChatGPT's "thoughts," you know, in air quotes, about something.

Leo: What are your thoughts?

Steve: What are your thoughts? They uploaded Samsung proprietary information in order to get ChatGPT to tell them what it thought.

Leo: Oh, lord.

Steve: And after three instances of that, Samsung...

Leo: Oh, god.

Steve: I know, has shut down. They imposed a 1K limit on the transactions with ChatGPT and said they are looking into measures for bringing a ChatGPT facility into their internal corporate network so that you won't have to go on the public Internet in order to use it.

Leo: I wonder, though, if there's any evidence that ChatGPT does save that information.

Steve: I think not evidence. But, for example, just here was the point being made that, if an account is stolen, then the past queries are apparently available...

Leo: Yeah, it saves a certain number of past queries. Not all of them.

Steve: Right, through the API.

Leo: It resets it, yeah. This is what I got in Italy. [Italian]. And then this is actually from the ChatGPT folks. They posted an English one, as well, which is nice. But, yeah, I thought, well, I heard that. So I said, well, let me try to use it because I have an account. It said "No, we regret to inform you we've disabled ChatGPT for users in Italy."

Steve: Oh, that's right. I forgot that's where you just were.

Leo: I was in Italy. I tried to use it.

Steve: Because we covered the news that Italy was saying it represented a privacy threat, and so they were just not going to make it available. Sorry about that. We're just saying no.

Leo: It's the same question, though, is - and ChatGPT denies it. They say, well, you know, we don't scrape everything. We scrape stuff that's publicly available.

Steve: Right.

Leo: So if they don't save stuff, and they only scrape public stuff, I don't know if there really is...

Steve: It's just, yeah, in fact, that was the point we made on the podcast was it's really just a different kind of spider than Google.

Leo: Yeah. It's a spider.

Steve: Which is indexing and saving the entire web. It's just it has a conversational interface, instead of one that's just sort of, you know, put in some keywords and see what hits you.

Leo: Exactly. This stuff will shake out. I don't think it's, yeah, I think it's a short-term issue.

Steve: I think we should take our second break. Then we're going to talk about what ESET found in some decommissioned routers that they purchased. Whoopsie.

Leo: Whooooo. Reminds me of when Simson Garfinkel bought a bunch of hard drives on eBay, and they were un-erased hard drives from old ATM machines and had all sorts of bank account information and stuff like that.

Steve: ESET made a very interesting observation in their posting on Tuesday, titled "Discarded, not destroyed: Old routers reveal corporate secrets." Uh-huh. Get a load of what they wrote. They said: "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" - shock, I tell you - "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 available. We're 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 been cybercriminals establishing an entry point and then a foothold into networks. They then spend time and resources conducting sophisticated extraction of data, exploring methods to circumvent security measures, and then ultimately bring 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. And 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 [ESET] attempted to alert them to the issues 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 and wiping'" - whoops - "a process that had clearly not taken place," they wrote, "and others just ignored our repeated contact attempts." They said: "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."

They said: "We have published the details, except the companies' names and data that would make them identifiable, in a white paper. The white paper also contains some guidance on the processes that should be followed, including references to NIST special publication 800.88 Revision 1, '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."

So I've got in the show notes, I have a link to ESET's white paper which is titled "How I (could've) stolen your corporate secrets for $100." And I also have a link to NIST's special publication 800-88 Revision 1, "Guidelines for Media Sanitation."

Anyway, in that white paper they provided a summery breakdown, just some bullet points of what they found by percentage. They said 22% of the routers contained customer data. 33% exposed data allowing third-party connections to the network. Think about that. One in three of the routers they purchased for 100 bucks off eBay or somewhere 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% also contained router-to-router authentication keys. 100% of them contained one or more IPSec/VPN credentials, or hashed root passwords. And finally, 100% had sufficient data to reliably identify the router's former owner and operator.

So again, you know, wow. Just a heads-up to make sure, if you're in charge or know somebody who is in a sufficiently sized organization, I would not trust a third party. I mean, how difficult is it to reinitialize the configuration in a router? That's, you know, return it to factory settings. It's not difficult. So don't pull it out and toss it on a pile to some third party that in this case apparently wasn't doing anything. And Leo, to the point you were making, I'll note that through 2023, four months so far, I've been having some very similar experiences of my own. I've been purchasing specific old drives from eBay.

Leo: Oh, yeah, of course, for SpinRite, yeah.

Steve: Right, when a SpinRite tester has reported that they've had some weird behavior from a specific old drive. And as SpinRite is running, one of its more popular screens is one which flashes up snapshots of the data which it's obtaining from the drive it's scanning. You just see it flashing by, but it's kind of mesmerizing, the way we used to watch all the blocks on a defragmenter move around the screen. It's sort of like that.

Leo: Steve, I've spent many hours staring at that screen. I know exactly what you're talking about.

Steve: We all have, Leo. It's an embarrassing truth.

Leo: No, but that DynaStat, I love that stuff, man. That's cool, man. I love it.

Steve: Yeah, yeah. So NTFS file system metadata has a particular look to it which I've learned to recognize. And these drives are 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. Oh, there you go, Leo.

Leo: And we're defragging here.

Steve: We're defragging. Wow. Love it.

Leo: That's interesting. They didn't erase those, either, yeah. I can understand how you would think, oh, well, a drive I have to erase. But a router? What could that possibly contain?

Steve: Right. Turns out those are the crown jewels of accessing inside of a corporate network. Makes sense when you think about it, but just make sure that you do think about it.

Okay. While we're on the topic of routers, let's take a look at last week's report, again from the UK's NCSC, regarding "Jaguar Tooth," a Cisco...

Leo: That's a terrible name.

Steve: I know, Jaguar Tooth, a Cisco router targeted malware. And this serves as a perfect case study. Jaguar Tooth is a system of backdoor Trojan malware developed via exploitation of a long-since-patched SNMP vulnerability. It's CVE-2017-6742. So there are two things here. First of all, 2017 tells you that it's now, what, six years ago, 2017. And 6742 reminds us of those quaint days six years ago when CVEs had four-digit numbering for their individual CVEs. We have to use scientific notation these days. This vulnerability was first announced by Cisco on the 29th of June in 2017, when updated and repaired software was made available by them. Cisco's published advisory included details of workarounds, including some of limiting access to SNMP from trusted hosts - imagine that - or by disabling several vulnerable SNMP API branches, which are known as MIBs.

Okay. So this amounts to another of those issues I so often have about policies versus mistakes. Backing up a little bit, SNMP is the Simple Network Management Protocol. It's 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. That's just nuts. 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 seen by 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. You know, it's just ridiculous. Let alone outsiders, keeping outsiders out. If external SNMP packets cannot reach the SNMP

service, then vulnerabilities in that service will never become an issue in the first place; right? So even if there's a mistake, if your policy is to firewall the SNMP service so that it isn't available, then no problem.

So once again, policies versus mistakes. Mistakes happen by mistake; okay? But policies happen by policy. In other words, on purpose. And mistakes don't need forgiveness, but policies don't deserve any. And 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 in SNMP from six years ago enables a stack-based buffer to be overflowed - whoever heard of such a thing - enabling control of the instruction pointer which can be used to gain remote code execution. This exploit uses Return Oriented Programming, the idea there is that until you're able to get your own code running, you need to use the code that's already there. And it turns out that an operating system like Cisco's IOS is full of subroutines. At the end of a subroutine is a return instruction which returns to where it came from, from where the subroutine was called. And first bad guys and then good guys figured out that the last few instructions just before the return can be useful. They can load something in a register. They can add something to a register. They can do useful things.

So clever hackers who have access to the same operating system as the one that they're attacking look at all the return instructions, and all of the instructions just in front of them, that are just before a return. And they're able to knit together the execution of code they want by jumping, deliberately jumping to just before subroutine's end, do a couple things, and then return to them, and then do it again, and again, and again, and again. Using just little snippets at the end of all the subroutines that exist in order to get done what they want done, which in this case is to incrementally assemble their malware in RAM. Again, it's a little frightening how sophisticated these attacks have been.

Turns out the vulnerable function targeted by this exploit is reached using the SNMP Object Identifier, the so-called OID, which corresponds to alpsRemPeerConnLocalPort. By appending additional bytes to the end of the OID, a stack-based buffer can be overflowed, which tells us that the expected length of this OID is how much space was created on the stack to hold it. But they didn't check to see if the OID was actually that long. They just went ahead and parsed it. So you literally - an OID is a weird-looking thing in SNMP. It's 1.3.6.2.27.14 dot.

So the idea is each of those dots is a branch in a tree. And so SNMP represents all of the little settings that you might have in a router by the end points of this incredibly richly branching tree where you address it by following this crazy dotted syntax all the way out to the end. Well, it turns out you could just keep adding things with dots in this broken version of IOS from 2017 and put a bunch of stuff on the stack, which then you can cleverly design what's there in order to go out and execute little code snippets for you.

One of the side effects of this vulnerability is that any ASCII characters in that additional OID bytes are converted to uppercase, which can be inconvenient, so the attackers get around that. 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 RAM, 4 bytes at a time. Once the Jaguar Tooth payloads have been copied into memory, they're individually executed by overflowing the return address of the vulnerable function with their location in memory.

Once Jaguar Tooth is running, it uses TFTP, that's the Trivial File Transfer Protocol, like a really reduced subset of FTP, to exfiltrate pretty much everything the router knows about all of the peers that touch it, its configuration, and all those things we were just talking about that you don't want to decommission, you don't want to leave in your decommissioned router. The router's ARP table, for example, is dumped to obtain the MAC addresses and IPs of all the internal machines that have recently touched the router. And of course the bad guys now have a foothold in a border router, able to run whatever they choose. From there on, it's going to be pretty much bad news.

So as I noted earlier, Cisco responsibly updated their IOS, that's what they call it, the Internetwork Operating System, back in June of 2019. The bad guys probably assumed that there would be routers, correctly assumed there would be routers that had not been updated in six years and were still running this IOS software from back then. So sure enough, 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. So another example of the problem that we have in our industry that we have not been able to figure out. We are unable to write software that doesn't have these kinds of problems, no matter how much we try and how much focus and attention we give to it; nor are we able to essentially recover all of the software that's already out there with known problems which nobody is taking the measures to fix.

Okay. Time for some happy news.

Leo: Well, it's about time. I've been waiting for this for 18 years.

Steve: We may have some really happy news at the end of this.

Leo: Oh. Oh, good, okay.

Steve: But first some interim 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, SecurityResearchLegalDefenseFund.org with no spaces or hyphens or anything. They explain themselves in one long line. They said: "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." Which is very cool.

So they break this down in three statements. The first is their mission statement: "The Security Research Legal 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." So, yay, you know, this is great news.

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 legal 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 guarantees 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." Okay, duh. "Also 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. It 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, finally, Board approval."

So I think this is great. You know, through the years here we've talked about this problem where well meaning, typically amateur hackers who are not backed by an organization attempt to inform an organization of some significant problem that they've stumbled upon and identified, only to have the organization's management freak out and aim their law enforcement and their attorneys at the hapless hacker.

Leo: That's what happened to Randal Schwartz at Intel, and he was arrested and prosecuted.

Steve: Yes. Yes. I mean, it's so wrong. So it seems like, you know, this would be a terrific backstop for such situations. I did a little bit of more research and legwork. SC Magazine knew a little bit more about this. It turns out that Google is like the main anchor for this. They said: "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," and they said, "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."

Anyway, they go on at some length. The council, that is, this Board of Directors, 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. So anyway, I just think this is great news, that there will be this sort of a formal legal defense fund that will be backed by people who understand the nature of the business and will be there to support people who get themselves in trouble when they're absolutely not being black hat, not attacking, not doing anything but trying to let

someone know that they've got a problem, and end up being attacked and sued in the process as a result.

Okay. So a couple quick updates. Last week Firefox users moved to v112.0.1 to fix exactly one problem, which was like, I don't think I've ever seen that happen. But because it seemed to be of great concern, Mozilla immediately moved on it. 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." Now, that can happen, right, because cookies are meant to expire. And so if the cookie dates were set into the far future, which shouldn't happen, that could be a problem, and it would be a big problem for people who expected to be, you know, to remain signed in. I didn't notice any problem. And I am once again, after using Google and Bing browsers for a spell, I'm happily back at Firefox, Leo, where I assume you still are.

Leo: Never moved. Never moved. So I've tried, but I never moved.

Steve: I know, yeah.

Leo: They kept giving me, they said, you should try this or that or Opera or Brave or...

Steve: I've been watching Paul. He's been, like, roaming all over the place.

Leo: He goes all over. Yeah, Firefox. You know why? Because I want to support an ecosystem that has more than one browser engine. Everything else is Chrome, or Chromium.

Steve: Yes. And Paul is a little reminiscent of Jerry, who used to say, "I do all these dumb things so you don't have to."

Leo: Yes. You're talking about, of course, the great Jerry Pournelle, who is much missed, I have to say.

Steve: Yeah. He was great.

Leo: He was a great guy.

Steve: Also you'll be glad to hear this. Kubernetes received a security audit. The NCC Group concluded and published a new security audit of the Kubernetes automation platform. Nothing significant found.

Leo: That's huge because it's very widely used yeah.

Steve: Yes. That's very cool. And finally, Chrome fixed a zero-day. They released 112.0.5615.137 or 138, which fixes eight security flaws, including a new zero-day. And

I'm not sure I understand this, unless it's a - oh, yeah. I'm sorry, I was looking at the number rather than the date. A zero-day, CVE-2023-2136. So it was discovered by Google's TAG team, so internally, although being a zero-day, they're not like Microsoft who calls it a zero-day if they learn of it by surprise. It's a zero-day when it's found being exploited in the wild. So anyway, this was patched last week.

Leo: And I wasn't aware of that. I did not know that. That's interesting.

Steve: Yeah.

Leo: Microsoft has a different definition.

Steve: Yes, they've got their own private definition, which is weird. It allows them to, I mean, you wouldn't think they would be declaring...

Leo: No, it makes more things zero-day than you would normally expect.

Steve: Exactly. Exactly. Which, you know...

Leo: They're not in the wild yet. It's still a zero-day? Wait a minute. Yeah.

Steve: Yeah. Google did note that it was under abuse by a surveillance vendor.

Leo: Oh.

Steve: So just exactly the type of slime that we were talking about before, somebody selling these things to third parties.

Leo: A data broker type?

Steve: Yeah, a surveillance - well, no, like an NSO Group.

Leo: Oh, geez. Oh, god.

Steve: It's a vendor selling surveillance capability.

Leo: Terrible.

Steve: Yeah. Okay. I'm very excited about what we've got coming next, an end-to-end encryption proposal. Let's do our final break.

Leo: Yeah. We'll do a little sponsor - yeah.

Steve: I'm going to share two open letters and then an idea that I had that might just work.

Leo: You've proposed things in the past. But boy, now more than ever we need it. EARN IT is back. The UK is about to do this. I mean, we need something. We need a solution that works for everyone. And if anybody could come up with it, it'd be you, Steve. So I look forward to hearing this. But meanwhile I just want to tell, first of all, I want to thank all of our Club TWiT members because you are the salt of the earth. You're the people who, you know, you could, and most of you do, like 95%, I think it's more like 98% of all of our listeners are very happy to listen for free. We've been doing it that way for 18 years. That's fine. Did you know, Steve, we had our 18th anniversary while we were gone, on April 17th? 18 years.

Steve: I know because you and I are in - yeah.

Leo: Right along in there, yeah. We were the second show; right?

Steve: Yup.

Leo: Yeah. So, and, you know, I'm happy to have you, and please listen for free as long as you want. But I really have to say appreciation to the listeners who say, you know, I like what you do, and we want to support it. We want to keep it on the air. We want you to grow. And they contribute. And it's not much. It's seven bucks a month. That's one, you know, half caff, half decaf cappuccino frappuccino with a twist. That's it. We really love our club members. They're the people who are most engaged, most involved, really support what we do. And thank you. Because of you, we're able to do more than ever before, and keep the lights on. TWiT.tv/clubtwit.

Now let's talk about Steve's proposal that's going to solve all our problems.

Steve: Okay. Maybe it's a good idea. Anyway, we'll see. 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 it's 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 individuals are suspected of illegal behavior or not.

Okay. So 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. We'll see.

Leo: Let's submit it for their approval, yes.

Steve: I'll share it with everybody, and we'll see what you think. But okay. First things first. Last week's news is that the CEOs 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 - they make some really great points - I want to share what the heads of today's secure messaging companies just wrote to the UK. They addressed it to anyone who cares about safety and privacy on the Internet. Okay, a little bit loaded there, but okay.

They said: "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" - that's 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 of the bill 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 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?

Leo: Yeah, it doesn't matter.

Steve: How are we supposed to interpret it? We're not interpreting it, we're just reading it.

Leo: You shouldn't interpret that, you know.

Steve: Wow.

Leo: It's, yeah.

Steve: And then they said: "Global providers of end-to-end encrypted products and services cannot weaken the security of their products and services to suit individual governments."

Leo: There you go.

Steve: "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." And it was signed by, they said, "those who care about keeping our conversations secure." And it was the CEOs of the companies I first mentioned.

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 workable. Now we have an 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. Now, 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 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: We have Australian Hi-Tech Crime Centre and the Australian Federal Police; Child Exploitation and Online Protection Centre in the UK; Colombian National Police; the Cybercrime Coordination Unit of Switzerland; the Dutch National Police; Europol; Interpol; the Italian Postal and Communication Police Service; the Korean National Police Agency; the Royal Canadian Mounted Police; the New Zealand Police; the Ministry of Interior for the United Arab Emirates; the Philippine National Police; and U.S. Immigration and Customs Enforcement (ICE), which is a division of the DHS."

Okay. So this group of actual police forces, law enforcement, have collectively authored and sent their own open letter, this time to Meta, asking Meta 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, you know, Child Sexual Abuse Material. So here's what they said. This is their open letter.

"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 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 organizations to tackle the threat of child sexual abuse. The VGT issued its first position statement on end-to-end encryption in 2021. This statement highlighted the devastating impact end-to-end encryption 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 end-to-end encryption 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 end-to-end encryption, at scale alongside robust safety systems that maintain or increase child safety. And where the child user base and risk is high, a proportionate investment and implementation of technically feasible safety solutions is paramount." They said: "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 child sexual abuse 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 interfere with the effectiveness of such safety systems threaten to undermine these successes which have been consistently built upon over previous decades."

Finally: "The announced implementation of end-to-end encryption 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 end-to-end encryption will effectively match or improve their current detection methods."

Okay. So everybody's gearing up here and staking out their positions. It feels like it's coming to a head. 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 from their positions, especially not those like Telegram, Signal, and Threema, whose entire purpose is providing secure end-to-end encryption. As we know, 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 the 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 should, you know, 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. Why don't we arrange to only monitor children? The pending laws and legislation would be 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 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 within 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.

So think about 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 for them 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 and set up 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 ever monitored. The online slimeballs who seek to take advantage of youthful trust and innocence know that all of their communications with an underage target is 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, set up 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, et cetera. 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 the 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 occurred to me yesterday, and I can't find much fault to 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 among them would be chosen so that the youngest user remains protected, and all remain in compliance with local laws.

In all 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 middle ground 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 want their services to remain legal where monitoring of minors is required by law.

Leo: Okay. I don't think they're actually proposing monitoring kids. They want to monitor adults for trafficking CSAM.

Steve: Well, oh. So, okay, I guess I missed that.

Leo: I think you missed the point.

Steve: So I thought it was to...

Leo: It's to protect children.

Steve: ...keep it from children.

Leo: No. Well, that's, I mean, that was a minor part of what Apple was proposing. But that's not what all these other things are about. They're about, whether this is true or not, but...

Steve: I guess I'm too far out of the loop, Leo. I really don't...

Leo: The thing they're trying to prevent...

Steve: I don't understand what the bad guys are doing.

Leo: The thing they're trying to prevent is adults from trafficking in child sexual abuse material. That CSAM database Apple was going to put on your phone comes from NCMEC, the National Center for Missing and Exploited Children, is a database of child pornography that adults are trafficking in.

Steve: That's really creepy.

Leo: Oh, Steve, you are very innocent. I did not realize you didn't know about this. Yeah, no, that's the whole - of course it's horrible. It's horrific. But it's not about monitoring children. I mean, that's part of it, I guess. You know, Apple has a thing where...

Steve: I guess I was focused on the whole grooming aspect of it. So, wow.

Leo: Yeah, that's not the issue. The issue is they really want us - the problem is all of these cloud services are storing tons of data. They don't want them to be storing child sexual abuse material. That's what CSAM is. It's images, graphic child porn images.

Steve: It's hard to even imagine that.

Leo: It's not kids sharing these, it's adults sharing these. That's the people they're going after.

Steve: What was it on "Saturday Night Live" where at the end of something, someone said, "Never mind"?

Leo: Roseanne Roseannadanna. No, there is a part of this, you're absolutely right. And Apple, by the way, is doing this. So if a child receives a naked picture - the parents could turn this on - or wants to send - this is actually what they're really trying to stop is children sexting each other - or wants to send a naked picture of themselves, Apple will say, "Oh, you don't want to do that," to the kid. And the parent could turn on the thing where it will warn the parent that this is going on. That's...

Steve: We have talked about this before. I didn't realize that that was actually in place, that that is in place now?

Leo: That's in place now, yeah. They turned that on. But that's not what all of this - that's not what the FBI or the UK wants. What they want is to break encryption for all of us so that pedophiles can't communicate.

Steve: Oh, Leo. It's hard to even picture that happening.

Leo: Oh, it's horrible. And that's why, by the way, it's a very useful tool for law enforcement because nobody's in favor of that, and you don't want to say, well, don't break encryption because they'll say, well, are you in favor of child porn? No. But the problem is you can't break encryption for pedophiles and still leave it intact for us.

Steve: Yeah, I certainly understand the mistake I made.

Leo: Well, you protected the kids, which I agree is a good idea. You've proposed in the past, quite a few years ago, an escrow system which I think still might not be a bad idea, third-party escrow system, like Apple holds the keys. But most of the people on that letter, for instance, just don't want any backdoors at all, or any escrow system, or any keys. These should be end-to-end means only you and the recipient have the key, no one else. And the problem is, what law enforcement says, is well then pedophiles can exchange all of this stuff, and we won't know.

Steve: Wow.

Leo: Yeah, I know. I'm sorry, I didn't mean to burst your innocence, Steve. I want to pat you on the head.

Steve: Wow.

Leo: Yeah, there's some bad people out there. There's some really bad people out there, more than we think, I think. But nevertheless.

Steve: Well, okay.

Leo: That's why it's a complicated thing; right?

Steve: I don't have any answer.

Leo: Yeah. That's why it's a really complicated thing.

Steve: I mean, it's hard to embrace the problem.

Leo: Yeah, I know.

Steve: Ugh.

Leo: Oh, Steve, I'm so sorry. I didn't know you didn't know. Go talk to Lorrie; would you? Ask her.

Steve: Oh, god.

Leo: Ask her about it. Steve, thank you for a wonderful show. You have a solution for that one area of it; right. And in fact that's basically what Apple's doing, by the way.

Steve: Well, and I was going to share these two open letters because what we are seeing is both sides escalating this thing.

Leo: Right. Oh, they're at loggerheads, and that's why it's tough. It's really tough because nobody wants to say, oh, yeah, we want to facilitate child porn. But it's much more than that. And this is, you know, Phil Zimmermann, creator of PGP, I did a Triangulation...

Steve: I think I'm in favor of monitoring, Leo. This is just so horrible.

Leo: Well, that's one of the reasons why that's what they use. That's the specter that they raise because who's going to be in favor of that? But I personally don't think that's the only thing they want. I mean, look, law enforcement says, if we could see everything going on, there would be no crime. We could stop crime, nip crime in the bud. The Constitution says, yeah, that's true; but people deserve - they have a right to privacy in their own home.

Steve: Yeah.

Leo: So this is the tension. This is why it's very difficult. I interviewed Phil Zimmermann, and we raised - this has always been the issue with - people said with PGP, look, that lets criminals exchange information freely. And Phil said, look, don't fool yourself. Law enforcement will say to you, oh, we're going dark. In fact, the FBI put out that paper, "Going Dark." They're not going dark. Technology's given them far more means of surveillance.

Steve: And what you're referring to when you talked about the thing I proposed a long time ago was the idea, a means of coming up with the equivalent of a search warrant.

Leo: Right.

Steve: Where under warrant a specific individual's communications could be monitored appropriately.

Leo: Right.

Steve: And again, though, we do have the problem of it not getting out of control, of it not being abused.

Leo: And that's the premise that WhatsApp and Signal and everybody else says, if you've got a backdoor, it will leak out. We know that historically. It will leak out. And so you can't have any backdoors. It's not safe. It's, look, Steve, I sympathize with your feelings because it's a horrible thing. Nobody's supporting pedophiles. But again, law enforcement, yeah, we could give them the right to see everything. That's what they'd like. But that would also mean that no one would have any privacy at all. It would eliminate crime, but at what price?

I'll let you think about this for a while, Steve, and we'll come back next week. Steve Gibson is the man. GRC.com is the site, the Gibson Research Corporation. Go there, get SpinRite.

Steve: I'm going to go back to work on SpinRite.

Leo: Yeah, stick with SpinRite. This is why you didn't want to do an - you were working on an encryption tool. And you knew...

Steve: Yeah, CryptoLink because, yeah, I knew the government was going to say...

Leo: They would come for you.

Steve: They'd be unhappy with absolute encryption.

Leo: Yeah. GRC.com has this show.