SECURITY NOW!

Transcript of Episode #921

OSB OMG and Other News!

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

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-921.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-921-lq.mp3

SHOW TEASE: it's time for Security Now!. Steve Gibson is here with some really interesting topics. Can the UK force companies to break their encryption? And what is to be done about it? We'll also talk about an unusual case, but some evidence that maybe it isn't always a good idea to rush those security updates out. And a word, if you don't mind, from the creator of cURL. It's all coming up next, plus a lot more, with Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 921, recorded Tuesday, May 2nd, 2023: OSB OMG and Other News.

It's time for Security Now!, the show where we cover the latest news in security. The show that gets longer every week because there's never any lack. Mr. Steve Gibson. Good day, sir.

Steve Gibson: It is funny. Sometimes I have an occasion to go back in time and look at some of our earlier shows.

Leo: The 20-minute versions, yeah.

Steve: Yeah. And I think, how did we get anything done in 20 minutes? You know, we're still busy saying hi to each other. And by the way, Leo, have you run across "Barry" on HBO?

Leo: You know, I've watched the first three seasons, and I saw that a fourth has arrived. So, yeah, we'll be watching it. You like it.

Steve: Yes. It starts a little slow. We're not sure. And what was weird was I was used to seeing the name Alec Berg, and I couldn't place it. So I googled him. Well, he was the guy behind "Silicon Valley."

Leo: Right. A talented writer, yeah.

Steve: And it does have - yeah.

Leo: And Bill Hader's hysterical. So I love Bill Hader.

Steve: What a good job.

Leo: And is the bald guy in it still?

Steve: Yeah. He's in it all the way through.

Leo: I love him.

Steve: And they are, you know, they're making him increasingly gay, so they're having fun with all of that.

Leo: He's such a great character. He's, what, a Chechen mobster, but sensitive. He's a sensitive...

Steve: He has style.

Leo: He has a lot of style.

Steve: A lot of style, yeah.

Leo: Yeah, I really, I quite enjoy that.

Steve: They really blinged him out in one of the episodes in the fourth season.

Leo: Did they really? Oh, that's great.

Steve: Yeah. And the way I found it was that somehow I clicked - I saw something that I think it was on Vox or something that said the show on HBO that you're not watching just started its fifth and final season, or is about to start its fifth and final season. And I thought, oh, that's interesting. We're ready for something. Because we're rewatching "House" because we're sort of out of things. So anyway.

Leo: Well, you and I after - and every once in a while I'll get a text from Steve saying "You've got to watch this." You and I after maybe we'll compare notes because we've found a few shows that we're quite enjoying.

Steve: Ah.

Leo: I know you're a "Succession" fan, so you've got to be watching this on Sunday nights.

Steve: Yes, well, actually, we're holding off.

Leo: Oh, that's smart.

Steve: So that we can do the whole thing. Because these cliffhangers...

Leo: But it ain't easy because there's a lot of spoilers on the Internet, and there's some real spoilage that can happen. So just close your eyes.

Steve: Well, the good news, Leo, is I don't use the Internet.

Leo: Oh, smart. Yeah, stay off Twitter, and you're okay.

Steve: Yes. I do not think...

Leo: I do not have that luxury. I wish I did.

Steve: I don't think WiFi works, and I don't use the Internet.

Leo: You know what's funny, yesterday I was telling Lisa, because we've been having IoT troubles and all sorts of stuff, and I said, "I can't wait till I can retire and I don't have to use technology anymore. We're going to move to a cabin in the woods. All I need is hot and cold running water and electricity. Forget the Internet." She doesn't believe me.

Steve: So there, Leo, is half of our show from the old days. That consumed 10 minutes.

Leo: Right there. The happy talk about - yeah.

Steve: None of our listeners will ever be able to get back. Okay. So this week, because the UK's Online Safety Bill continues to stir up a hornet's nest of worries and concerns within many industries, we're going to examine WhatsApp's reaction to Signal's "We plan to walk" position.

Leo: Oh, yeah. Oh, yeah.

Steve: And Wikipedia's concerns over the bill's age verification requirements. And, undaunted, I have another idea that might be useful.

Leo: Never stop coming up with ideas, Steve. We love you for that.

Steve: That's right. So we also have a new UDP reflection attack vector, lots of amplification in it; a welcome, and late, update to Google Authenticator; more NSO Group client news; a Russian OS?; the unintended consequences of releasing updates for routers that won't actually ever be updated; a smart move by Intel with pre-release security auditing; yet another side-channel attack on Intel CPUs; cURL's maintainer implores Windows users not to delete it; and VirusTotal gets AI. So I think a great podcast. I titled this "OSB OMG and Some Other News."

Leo: Well, we'll get to it PDQ. The Picture of the Week, Steve.

Steve: Okay, so...

Leo: This is wild.

Steve: And Leo, it has to have been photoshopped. Don't you think?

Leo: Yeah. Nobody would do this on purpose.

Steve: No, no one could build this. So for those who don't have the advantage of video, what we have is a bizarre swing set where just past the bar that holds the chains mounting the swings is a brick wall so that, you know, you can't swing because the moment you would go past center you'd be hitting this brick wall. Literally. Maybe kids could sit there, like push off of the brick wall with their feet? I mean, I just don't know what is...

Leo: But as soon as it came back, boom, you're going to hit the wall.

Steve: It's not good.

Leo: That's crazy. That's crazy.

Steve: Yeah. So but the picture captivated me, so it made it into this week's show.

Leo: I believe it, Steve. I think it really happened.

Steve: Do you?

Leo: No. I'm just trying to give you credibility. No, that's definitely real. What I like is it's like the best built swing set you ever saw. I mean, this thing is...

Steve: It's very sturdy.

Leo: Sturdy.

Steve: Except it doesn't do anything good. Anyway, so it came with the caption that I used, which is "Security should never be added as an afterthought."

Leo: Love it.

Steve: As if to say that the brick wall was added after the fact and, whoops, you know, it broke the swings. But I'll bet the swings wouldn't work without the brick wall. So, you know, anyway, I just - it was wacky enough that I thought, okay, we don't really have an explanation for this, but we'll just put it in the show notes.

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

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

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

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

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

Okay. So this leaves us with a question. What's going to happen? No one knows, which is I think what makes this so interesting. There doesn't appear to be any possible way to compromise. Either governments have no officially sanctioned way to monitor communications, or they do. It's glaringly binary. If the UK's current proposal were to be enacted into law, there would presumably be some length of time provided for encrypted service providers to come into compliance.

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

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

JD Supra, a legal news site, had this to say about the Online Safety Bill's penalty provisions. They wrote: "Penalties available include fines of up to 18 million pounds or, if higher, 10% of global turnover" - meaning their global revenue - "and which for the larger providers could be significant sums. In addition, it can impose business interruption measures including, ultimately, service restrictions.

"A particularly controversial measure has been the availability of up to two years' imprisonment for senior managers who suppress, destroy, or alter information requested by Ofcom" - remember that's the UK's communications regulator - "who fail to comply with, obstruct, or delay Ofcom when exercising its powers of entry, audit, and inspection; for providing false information; or for employees who fail to attend or provide false information at an interview. A recent amendment also provides a further offense where a senior manager has 'consented or connived in ignoring enforceable requirements, risking serious harm to children.'

"For these purposes, a 'senior manager' is if the individual plays a 'significant role in the making of decisions about how the entity's relevant activities are to be managed or organized; or, B, the actual managing or organizing of the entity's relevant activities." And they finish: "The OSB is in the latter stages of the legislative process; and while substantive amendments may still be made, it is likely to receive Royal Assent by midyear. The OSB may also have the unintended effect of causing terms of service to be watered down as to what content a service may contain. Ultimately, some providers may decide it is simply too difficult to comply with, and instead block UK users."

Okay, well, there's obviously no way Tim Cook or Mark Zuckerberg are going to prison over this.

Leo: Yeah.

Steve: You know? Nor are they going to give up 10% of their respective companies' annual revenue.

Leo: No, that's like, for Apple, that's like $36 billion. That a massive amount of money.

Steve: Not going to happen. So ignoring the law doesn't appear to actually be a practical option. And note that the law was clearly crafted to make exactly that fact exceedingly clear. But I wanted to cover that case since it was theoretically one of the three possibilities.

Okay. So it turns out that I wasn't totally out in the weeds last week with my proposed design for a device-centric age verification solution, and something like that may actually be the right answer. Even if the Online Safety Bill (OSB) fails, as we all hope it will, in enforcing the monitoring of all private communications within the UK...

Leo: There is now a U.S. bill called the CSAM Act because they know, you know, nobody's going to vote in favor of child porn. So this is going to be a problem everywhere, not just in this case.

Steve: Yes, that was the thing that the Senate enacted on Wednesday that you're talking about?

Leo: Yeah, yeah.

Steve: Yeah. So, you know, regardless of how in this case the UK bill's authors want it to be "interpreted," you know, what the Bill requires is...

Leo: A back door. A back door.

Steve: ...side-channel monitoring, yes. Its extensive provisions, which require the moderation of all content made available specially to those under the age of 18, those provisions appear very likely to survive. So it appears that some answer to the unsolved challenge of online age verification is going to be needed.

Okay. So this was highlighted by some news that was covered four days ago by the BBC in their piece titled "Wikipedia will not perform Online Safety Bill age checks." Okay. So here's what the BBC wrote. They said: "The Wikipedia Foundation says it will not comply with any age checks required under the Online Safety Bill. Rebecca MacKinnon, of the Wikimedia Foundation, which supports the website, says it would 'violate our commitment to collect minimal data about readers and contributors.'

"A senior figure in Wikimedia UK fears the site could be blocked as a result. But the government says only services posing the highest risk to children will need age verification. Wikipedia has millions of articles in hundreds of languages, written and edited entirely by thousands of volunteers around the world. It is the eighth most-visited site in the UK, according to data from analytics company Similarweb."

Leo: Yeah. And of course many of those are kids. And it's just a matter of time before they decide that, oh, no, Wikipedia counts.

Steve: Right. Exactly.

Leo: And I'm sorry, yeah.

Steve: And that is the concern here. "The Online Safety Bill currently before Parliament places duties on tech firms to protect users from harmful or illegal content and is expected to come fully into force sometime in 2024.

"Neil Brown, a solicitor specializing in Internet and telecom law, says that under the bill, services likely to be accessed by children must have 'proportionate systems and processes'" - I know - "designed to prevent them from encountering harmful content. That could include age verification. Lucy Crompton-Reid, chief executive of Wikimedia UK, an independent charity affiliated with the foundation, warns some material on the site could trigger age verification. 'For example,' she said, 'educational text and images about sexuality could be misinterpreted as pornography.'

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

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

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

Leo: Ha. Ho.

Steve: Uh-huh, "...was an example of how community moderation can be effective." Meaning figure it out, and you guys do it. "He said the bill did not say that every service needed to have age checks, and it was expected that 'only services which pose the highest risk to children will be required to use age verification technologies.' Ms. Crompton-Reid told the BBC that while Lord Parkinson's remarks 'reassured' her, the charity did not want to be relying on" - as you said, Leo - "on future goodwill and interpretation of legislation. She said they would continue to urge that protections to community moderation were in the bill through measures such as the exception for public benefit websites like Wikipedia.

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

Leo: Oh. This just happened in Utah.

Steve: Uh-huh.

Leo: And Pornhub has withdrawn from the state.

Steve: Exactly. Yeah, exactly.

Leo: The problem is, if it's just England, you know, honestly, fine. Bye-bye. But it's not going to just be England, and that's the problem. It really is.

Steve: Well, and that's why I'm hoping that we're going to see, I mean, so this stuff is separate from encryption providers. I think encryption providers are going to just decide this is the hill they want to die on, and they're just going to say no.

Leo: That's what Signal says, yeah.

Steve: Yes.

Leo: Signal says no, we're not - it's not going to happen.

Steve: So last month...

Leo: And I think honestly encryption is going to end up being civil disobedience. Using encryption's going to be a form of civil disobedience. The problem is that's for individuals. It doesn't solve the problem for these big companies. I'm not sure how they solve this.

Steve: Right, right. So last month the Guardian published what they called an "Explainer" about the age-related aspects of the impending Online Security Bill. Their piece was titled "Will UK's online safety bill protect children from adult material?" And their subheading was sort of the crux of it, saying "Legislation puts duty of care on tech firms to protect under-18s, but does not mandate use of specific age-checking technology."

Okay. So they said: "The online safety bill is due to become law this year, and it imposes a duty of care on tech companies to protect children from harmful content. However, there are calls from campaigners and peers to toughen the legislation's provisions regarding pornography. Here's what the act proposes to do on adult material. The bill requires all pornography websites, such as Pornhub, to ensure children do not encounter their content. This will require age-checking measures. The legislation refers to stringent age verification - checking a user's age via government ID or an authoritative data source such as a person's bank - as a means of doing so. Breaches of the act carry the threat of a fine of up to 10% of a company's global turnover or, in extreme cases, blocking a website altogether."

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

"The primary regulator of legal pornography offline is the British Board of Film Classification, which gives pornography age ratings - R18 for the most extreme but legal content, or 18 - but it has no control over online content. Ofcom, the communications watchdog, already has the power to regulate UK-based 'video-sharing platforms' such as TikTok, Snapchat, and OnlyFans. These platforms are required to protect under-18s from videos containing R18 material such as pornography. The age-appropriate design code was introduced in 2021 and is designed to prevent websites and apps from misusing children's data. Under its terms, social media platforms would be breaching the code if their algorithms served adult material to under-18 year olds.

"Age verification has been a troublesome issue for the government. Age checking for pornography was announced as a Conservative policy in 2015. However, plans to introduce a nationwide age verification system for online pornography were abandoned four years later in 2019. The bill will not mandate use of specific technologies for age checking, although Ofcom will issue codes of practice on age assurance, which is the umbrella term for assessing the age of people online. Age verification is the term for the toughest measures, such as requiring proof of official ID.

"One solution is to use age verification companies that vet a user's age, via a range of methods including checking offline ID or bank statements, and then notify the porn provider that the person wishing to access their service, who is anonymized, is over 18 years old. Ofcom has said it will launch a consultation on protecting children from pornographic content, including on user-generated platforms such as OnlyFans, in the autumn. The government has indicated that there will be clear instructions to mainstream social media sites and search engines to prevent children accessing pornographic content on their services." That's right, search engines, of course, is another problem. "The bill requires sites to prevent children encountering what it terms 'primary priority content.' Because it qualifies as a 'user-to-user' service, subscription site OnlyFans is also covered by this part of the bill.

"We will not know what is primary priority content officially until it is defined in a statutory instrument that will be published after the bill becomes law. However, pornography is expected to be on that list, and it was listed as primary priority content by the previous culture secretary, Nadine Dorries, in a parliamentary statement last year. According to a timeline published by Ofcom, though, it could be more than 18 months after the bill is passed before these provisions come into effect. Social media sites and legal pornography sites will also be required to shield all users from illegal pornography such as obscene content and child sexual abuse material.

"The bill will update the law on sharing intimate images without someone's consent. In England and Wales there will be a new 'base offense,' where it is an offense to share an intimate image of a person if they do not consent, and the perpetrator does not believe they have consented. Currently, these offenses apply if the image is shared in order to cause humiliation or distress. The base offense will now apply regardless of the motivation, including sharing it as a joke, for social status, financial gain, or 'where there is no motivation at all.'"

Okay. So from all of that, the article's subheading seems to me the most pertinent to those of us who care and are interested in how these things work, and how they're done. The Guardian wrote: "Legislation puts duty of care on tech firms to protect under-18s, but does not mandate use of specific age-checking technology." So said another way, "We don't know how you're going to arrange to do what the new laws we have just written require you to do. But that's not our problem. It's yours because we said so."

Leo: This is because the old law required that you go to a pub, and I think there might have been like post offices, but somewhere where they had a system of checking ID, but most of the time it sounded like a pub, to verify your age. And everybody said, "You're going to send 15 year olds to a pub to verify their age?" And, oh, what about somebody who says "I'm an adult, and I want to view porn. I've got to go to a pub and say can I have a porn license, please?" I mean, this was a terrible plan, and they haven't solved it. They've just said, well, it's not our problem.

Steve: Yeah. Yeah, exactly. They're saying "We don't know how you're going to do it."

Leo: Just do it.

Steve: It's very much like the encryption problem; right? Well, we don't want you monitoring everyone's content, but you have to. You know?

Leo: I'm so depressed by all this, Steve. I really - it's just...

Steve: I know. I know. It's the collision that we've seen coming for a long time. So whether it's the perceived need to monitor everyone's communications all the time, in the off chance that something illegal might pass by, or the need to impose strict age restrictions on access to Internet content and behavior, it's apparent that the UK's legislators believe that they can ask for whatever they want, leaving it up to the tech companies to figure out how to do it, while all the while imposing penalties if they fail to achieve what might well be impossible, or at least impractical, like as you said, the previous legislation that they had passed.

So after last week's podcast I received a DM from someone who said: "Please don't go spending another seven years solving the age verification problem like you did with the online login authentication problem." To which I will formally respond: "Fear not. As they say, 'Fool me once.' I have quite thoroughly learned my lesson."

Leo: Oh, no, no.

Steve: And I'm having far too much fun working on new technology for SpinRite, where I can actually make a difference. So, you know...

Leo: You did propose some years ago, I'll have to find the episode, a kind of some sort of third-party key escrow.

Steve: Yeah, and that's still - that could still be done. But it doesn't solve - so that would allow privacy to be maintained and search warrants to be served. That is, so that was that. So the idea there was to try to bring into the end-to-end encryption space the similar U.S. constitutional protection where you need to prove to a court that you've done something wrong. On the other hand, now we see what the courts have been doing lately. So it's like, oh.

Leo: And also, as others pointed out, if you have a backdoor, those keys leak. I mean, it's hard to keep it. And you proposed a very good system, you know, you had - it was well thought out. I thought it was a very good idea. And it may be our last best hope because it may be that, you know, it's the better of two bad alternatives.

Steve: But it doesn't solve this problem that the UK wants to solve. They literally want to look to somehow screen all text messaging in case it might be grooming children.

Leo: Yeah. See...

Steve: And every image that you send in case it might be illegal content.

Leo: So it's not - it's everything. It's a phishing expedition.

Steve: Yes.

Leo: It's a broad net catching everything.

Steve: Yes. It is side-channel monitoring.

Leo: Yeah, it's terrible.

Steve: That's what this online safety bill requires. And that was the point at the beginning of this is that no one wants to provide any cover to those creeps to allow them to do what they're going to do.

Leo: No, of course not.

Steve: But they will do it using illegal technology if we open up the encryption.

Leo: Right.

Steve: And so that we're doing back-channel monitoring.

Leo: Right. It's not going to solve the problem, and it's going to compromise everybody else's privacy. And if you think a government only cares about CSAM and grooming, you're not paying attention.

Steve: Well, yes. You haven't been noticing where the cash for Pegasus is coming from.

Leo: Yeah, exactly.

Steve: It's coming from governments.

Leo: Exactly.

Steve: So in a little bit of happy news, recall that Signal's President Meredith Whittaker made some headlines when she told BBC News that Signal "would absolutely, 100% walk," and stop providing services in the UK, if required by the Online Safety Bill to weaken the privacy of its encrypted messaging system. That's the only stance that any entity like Signal, Threema, or Telegram could take; right? Because their entire existence is encrypted communications. But what's the position of the number one most popular and largest messaging app in the UK, WhatsApp? WhatsApp is used by more than seven in 10 adults who are online, according to the UK's communication regulator Ofcom.

Okay. So the BBC asked Will Cathcart, the head of WhatsApp. Will replied that WhatsApp would refuse to comply if asked to weaken the privacy of encrypted messages. Period. Full stop. He said that WhatsApp would rather be blocked in the UK than undermine its encrypted messaging system, if required to do so under the Online Safety Bill. He said: "We won't lower the security of WhatsApp. We have never done that, and we have accepted being blocked in other parts of the world." And he feared the UK would set an example, as you've said, Leo, other nations might follow.

Will added that undermining the privacy of WhatsApp messages in the UK would do so for all users. He said: "Our users all around the world want security. 98% of our users are outside the UK. They do not want us to lower the security of the product. We've recently been blocked in Iran, for example. We've never seen a liberal democracy do that. When a liberal democracy asks, 'Is it OK to scan everyone's private communication for illegal content?,' that emboldens countries around the world that have very different definitions of illegal content to propose the same thing."

After Will went on the record with WhatsApp's position, Signal's Meredith Whittaker tweeted, "Looking forward to working with @wcathcart and others to push back"; after which Will replied on Twitter, "And very important we work together, and honored to get to do so, to push back." So what appears to be forming is a bit of an insurrection, and this may be where the encrypted services companies decide they need to take a stand. Last week we saw their open letter to the UK regulators. So they all know each other, and they have each other's email addresses, and they're talking. That's all for the good.

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

Leo: There you go. Yeah.

Steve: So it would mimic the way iOS devices currently operate when messaging outside of Apple's closed and encrypted ecosystem to Android devices.

Leo: UK or Android. Both green.

Steve: Yup, exactly. And Leo, speaking of green, let's make a little green.

Leo: All right. We need some green right now. This is such a good subject. I have a feeling there's going to be mass civil disobedience. You know, I hope you're saving your crypto code, all of you, and coding what you need to have it because it's just not going to be okay. The problem is all these companies - eventually companies are going to have to give in. If it's just the UK, maybe they can write it off. If it becomes the EU, or the EU plus Australia, or the EU plus Australia plus the U.S., you know, the Five Eyes, the companies are going to comply in the long run. They have to. In which case it's going to be up to individuals to preserve their own privacy.

Steve: Yeah. And I have - this is where my next idea comes in. We'll get to it in a minute.

Leo: Can't wait. Coming up. Yeah, that's exciting. Yes.

Steve: Here comes my idea.

Leo: Good. Steve's idea.

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

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

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

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

And I had my own firsthand experience with people being unhappily surprised. Leo, you used to introduce me by mentioning that I discovered the first spyware and in the process coined that term.

Leo: Right.

Steve: What happened was I discovered that a freeware utility I had installed on my Windows machine - as I recall it was an early version of WinZip - wasn't as free as I was led to believe. It was an "ad-supported application," and so it brought along an advertising DLL from a third-party company named Aureate. What I discovered was that this Aureate spyware was inventorying my machine, monitoring my actions.

Leo: Oh, I remember this, yup.

Steve: Monitoring my actions and usage, and then phoning home without ever obtaining any permission from me. I had no idea it was there, and to say that I was unhappy when I found something communicating behind my back without my knowledge or permission would be an understatement. So I created "OptOut," the world's first spyware removal tool, a bit of freeware which successfully removed Aureate and several other early forms of stealth spyware.

The reason I bring this up is that the management at Aureate shared with me some of the way-over-the-top enraged and nearly psychotic emails they were receiving from PC users who were more than just a bit unhappy to have used OptOut to discover that their machines had also been infected. The Aureate people said that the ad-supported software packages which installed their spyware - I mean their adware - were supposed to explain the situation and obtain their users' consent. I asked why they didn't have their DLL present its own permission dialog. They didn't reply. Anyway, the name Aureate had been ruined by my crusade, so the company renamed itself to Radiate, and not long after ceased operations.

Leo: Yeah. Victory.

Steve: The whole concept was really never viable, and after this no freeware developer wanted anything to do with them.

Okay. So what does all this have to do with the UK's Online Safety Bill? What occurred to me was that, if encrypted applications were going to be required by law to arrange some sort of side-channel government-mandated monitoring, you know, eavesdropping and surveillance, they should make very clear to their user that that's what they are doing, and not repeat past mistakes of doing things that people would find objectionable if they were clearly informed of what was going on.

So the presence of state-mandated communications surveillance should be placed front and center for every UK resident, and anyone they communicate with. The top of any application that's being forced to break into its users' privacy in order to comply with the UK's Online Safety Bill should clearly display against a red background the message: "This communication is being monitored by your government."

I imagine that the presence of that notice at the top of any communicating application might provoke a reaction similar to what happened when the news of the Aureate spyware broke. The UK government will clearly wish the fact of this being hidden from its citizens' view. But it should be there to serve as a constant reminder of what the country's politicians have decided is in the best interests of their citizenry. And could you imagine, you know, a little red banner up at the top, "This communication is being monitored by your government." That'll provoke some change.

Okay. So we've got some Closing-the-Loop feedback from our listeners. TWS tweeted: "Hi, Steve. Are you still using the ZimaBoard you mentioned many episodes ago? Would you still recommend? With the shortages of Pi's I'm considering them. Thanks for making SpinRite and the SN podcast."

In a word, yes, yes, yes. Oh, okay, wait, that's three words. The ZimaBoard is the best thing I've found for my own work during the development of SpinRite. It is perfect. And while I haven't taken a show of hands in the Spinrite.dev newsgroup, I keep seeing people referring to their ZimaBoard in passing. So I know I'm not alone. But also for other uses, it is really a perfect little machine. It's been around long enough that there are now a ton of YouTube How-To videos covering pretty much anything you can think of. Just go to YouTube and put in Zima, Z-I-M-A-B-O-A-R-D, and you'll see.

And while working, as it happens, to assemble today's show this morning, I received an email from them announcing a Star Wars Day discount of 20%. Their email said: "May the 4th be with you." Apparently there's a little bit of lisp. And their emailing said that the sale runs for three days, from tomorrow, May 3rd, through Friday the 5th, offering, again, a 20% discount. So yeah, ZimaBoard.com. And if I didn't already own five of them, I'd be purchasing some more. It was a real find.

But remember, unlike the Raspberry Pi, which is ARM-based, the ZimaBoard is Intel-based. That makes it incredibly useful to me, but you'll want to be sure that whatever you want to do with it you can do with Intel-based software. And definitely check out YouTube. And I just wanted to make sure that everyone knew. We have a bunch of news we'll be covering after I deal with a couple of these little blurbs from feedback from our listeners.

David Schofield, yeah, he said: "Good morning, Steve. SN listener since 2007, SpinRite user since my ST-225s." Which of course were Seagate 20MB drives.

Leo: Oh, yeah, I remember them. I had them, too.

Steve: And he said: "Encouraging news about Microsoft rewriting parts of Windows in memory-safe Rust." So his note is a perfect segue for me mentioning a bit of news I had encountered. David's direct message to me linked to an article in The Register. Here's just the top of their piece. They said: "Microsoft is rewriting core Windows libraries in the Rust programming language, and the more memory-safe code is already reaching developers. David Weston, director of OS security for Windows" - he's got to have an interesting job - "announced the arrival of Rust in the operating system's kernel at BlueHat IL 2023 in Tel Aviv, Israel last month. He said: 'You will actually have Windows booting with Rust in the kernel in probably the next several weeks or months, which is really cool. The basic goal,' he continued here, 'was to convert some of these internal C++ data types into their Rust equivalents.'"

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

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

And of course, as we know, all of the evidence suggests that we're not really making any headway with simply trying to be more careful using powerful but unsafe legacy languages. Our software is growing more and more complex, and people make mistakes. The adoption of newer languages which prevent those mistakes from proving fatal to a system's security looks like the only way we're ever going to get to the point where we start removing more existing bugs than we are introducing new bugs as we go forward. So yay to Microsoft for making this move.

Frank S. tweeted: "Dear Steve. I have two kids, and I am very happy that they are under seven. This gives me and the market some time to find best, or better, practices for children using the Internet. As a parent, I want my children to be safe, both physically and online. However, I don't think that we can stop existing CSAM images that are already out there. What we can and should do is try to prevent new cases and victims. I believe it is up to the parents to take better care and guide our children when growing up. And I don't want the government to spy on them. As a parent, I would like more tools to keep my children safe online."

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

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

Leo: Well, and I sympathize with Frank and as every parent. But I have to say the problem is not people like Frank or his kids or their access to the Internet.

Steve: Right.

Leo: The problem is from parents unlike Frank who don't care, who exploit their children. Most abuse of children comes from relatives and people they know or their parents know. It is not from caring parents like Frank. And frankly, it's not from the Internet. And I think that all this focus on, oh, they're groomers, they're going to get you in an AOL chatroom, and then all of a sudden you're going to go out and do child porn, is really, I think, misdirecting it. And maybe it's easier for legislators to attack technology.

Steve: Well, and it's call a "straw man"; right?

Leo: It's a straw man.

Steve: Yup. Which is not to say that it can't or doesn't happen.

Leo: Right.

Steve: But that the purpose is not sincere.

Leo: No, the real danger is inside the house when it comes to children. It's relatives and people they know.

Steve: Yup.

Leo: Protect them from those people. Frank, you're doing a great job. I don't think you have to worry about a technological solution to protect your kids. If you're just paying attention, that's all you need to do. It isn't that hard.

Steve: Yes, be involved.

Leo: Be involved. But the problem is not there. I mean, there are a lot of parents who aren't involved, who don't care, or worse, participate. And, you know, we do need to catch those people.

Steve: Yeah.

Leo: But there are ways to do that. And all you're going to do is drive these things underground.

Steve: Yup.

Leo: That's not going to solve it either because, if you're motivated, you'll find a way to get it done. Anyway, I agree with you, Steve.

Steve: Yeah. So we've been focused upon the UK so far, but this is probably a good time to mention what you had said earlier, Leo, a new U.S. federal bill was announced and unveiled last Wednesday aiming to regulate access by age to social media platforms in the U.S. Here's a bit of CNN's coverage of this new proposal.

They said: "A new federal bill unveiled Wednesday would establish a national minimum age for social media use and require tech companies to get parents' consent before creating accounts for teens, reflecting a growing trend at all levels of government to restrict how Facebook, Instagram, TikTok, and other platforms engage with young users. The proposed legislation by a bipartisan group of U.S. senators aims to address what policymakers, mental health advocates, and critics of tech platforms say is a mental health crisis fueled by social media.

"Under the bill, known as the Protecting Kids on Social Media Act" - that's PKSMA, okay, that doesn't say anything, that's good - "social media platforms would be barred from letting kids below the age of 13 create accounts or interact with other users, though children would still be permitted to view content without logging into an account, according to draft text of the legislation.

"Tech platforms covered by the legislation would also have to obtain a parent or guardian's consent before creating new accounts for users under the age of 18. The companies would be banned from using teens' personal information to target them with content or advertising, though they could still provide limited targeted recommendations to teens by relying on other contextual cues. It's the latest step by lawmakers to develop age limitations for tech platforms after similar bills became law this year in states such as Arkansas and Utah. But the legislation could also trigger a broader debate, and possible future court challenges, raising questions about the privacy and constitutional rights of young Americans.

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

And I'll just note that this discovery is not unique to social media platforms appealing to children. Precisely the same observation has been made by cable news outlets about how to engage and enrage their audiences to encourage viewership. I'm not sure what "gen" we're on now, Leo, X, Y, or Z. I've lost track.

Leo: Me, too.

Steve: But young people have grown up with the Internet and hundreds of cable channels, each with their own agenda. I hope they can figure out how to handle the mess we have made of this.

So speaking of messes, and then we're going to get to the week's news, Simon Zerafa, a good friend of the show, he tweeted: "Shodan. Twitter canceled our API access, which broke the ability to log into Shodan via Twitter using single sign-on. Email us at support@shodan.io if you're currently logging in via Twitter and would like to migrate to a regular Shodan account instead of using single sign-on." Yikes.

And this reminds us that aside from the well-appreciated privacy implications of using OAuth-style "Sign in with" whatever, Facebook, Google and so forth, authentication, if that third-party authentication service should ever become unavailable for any reason, you're hosed since the only way you're known by the site you're wishing to log into is courtesy of that other third-party entity which may no longer exist. So, yeah, another downside of the super convenient logon using somebody else. You'd better hope those "else" people stick around.

Leo: I had a little tiff with Dave Winer a few months ago because all of his online services, which are very cool, use Twitter OAuth. And I said, "Dave, can you just do something else?" He said, "No, no, this is good." I wonder how he feels now.

Steve: Yeah. All the people that were using that for free.

Leo: Not free.

Steve: Yup, exactly, not free. Let's take our last break, Leo, and then we're going to plow into the News of the Week.

Leo: Yeah. His defense was it's easy.

Steve: Uh-huh.

Leo: I'm not a fan, as you know, of third-party OAuth.

Steve: No. So we've previously talked about UDP reflection attacks. Unlike TCP connections, which are inherently bi-directional and therefore require packet round trips between the endpoints to establish byte numbering and other connection parameters, the UDP protocol is often referred to as "connectionless" because, although it's still possible to establish connections by mutual agreement, UDP doesn't have that baked into its protocol. This makes UDP the perfect protocol for DDoS bandwidth flooding attacks since the sender of a UDP query can spoof the UDP packets that they're sending out, spoof its source IP so that the recipient of a query will redirect its reply to the victim of the attack.

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

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

Because of the protocol's huge potential for DDoS attacks, both Cloudflare and NetScout have said that they expect the prevalence of SLP-based DDoS attacks to rise significantly in the coming weeks, once threat actors learn to exploit it. The only good news is that since SLP is transported over port 427, and since it has no business being exposed on the public Internet, like it's for printers and things to find each other, it's a way for a printer to broadcast its existence and be found, so it would only make sense in a LAN context. It's only exposed publicly due to the typical mistakes of having ports exposed that shouldn't be.

Anyway, because it has no business being out on the public Internet, I would expect that many savvier carriers like Cloudflare will already be proactively blocking that port traffic at their borders. They just don't need to allow port 427 stuff to get even close to its targeted victims. There's no reason for it. But that won't help any unprotected targets. So DDoSers have had another arrow added to their quiver. And of course they don't lack for arrows, unfortunately.

Google Authenticator, which was first released 13 years ago in 2010, has just been updated with an extremely useful new feature which I've had in my favored iOS OTP Auth app from the start: Cloud backup. In their announcement of this groundbreaking technology, they wrote: "One major piece of feedback we've heard from users over the years was the complexity in dealing with lost or stolen devices that had Google Authenticator installed. Since one-time codes in Authenticator were only stored on a single device, a loss of that device meant that users lost their ability to sign into any service on which they'd set up two-factor authentication using Authenticator." Very much like using Twitter to log on.

"With this update we're rolling out a solution to this problem, making one-time codes more durable" - that's what you want in your one-time codes, some durability - "by storing them safely in users' Google Accounts." What a concept. "This change means users are better protected from lockout, and that services can rely on users retaining access, increasing both convenience and security." Talk about upselling a small feature. Anyway, that's great. Somewhat odd that it took them this long to get it, but I wanted to make sure that everyone listening who might still be using Google Authenticator - I was at one point, before I moved over to OTP Auth - would know of this critically useful new feature. And I would imagine you definitely want all of the private secret keys that are in Authenticator to be backed up so that should you need them somewhere else, you can get them. And as you know, I'm...

Leo: It's not currently end-to-end encrypted. When are they going to do that?

Steve: What is not? Google...

Leo: The backup.

Steve: Oh, no kidding.

Leo: Yeah.

Steve: So maybe that's what they added was some means of having a secret key.

Leo: So it syncs, but it syncs unencrypted. So Mysk, M-Y-S-K, was the group or person who discovered this. And Google has confirmed it. And they said, well, we're going to add that later. But we thought it would be just faster just to put this out right now.

Steve: Because, Leo, after 13 years, we suddenly need to hurry.

Leo: Yeah, why not?

Steve: That's right.

Leo: So you haven't seen anything that - so data syncs between devices with the new Google Authenticator update could be viewed by third parties. Google says the app works as planned. Christian Brand of the Google Group Product Manager Identity and Security tweeted that this is our intention. Because I guess E2E would be hard. I don't know. The lack of end-to-end encryption also means Google has a transparent view into what services each account owner uses as it's being transmitted.

Steve: Ah, that's true. It's not client-side encrypted. So they're getting that.

Leo: Mysk found the app does not expose the 2FA credentials associated with the user's Google account. So that's still secure.

Steve: Okay, everybody, let me tell you about OTP Auth.

Leo: Tell us what the name is of the one you use. OTP Auth?

Steve: OTP space Auth, and the logo is a simple gray padlock. Very modest logo. And it does all of this correctly.

Leo: Yeah. And I've been using one that's open source called 2FAS Authenticator. And the way it works is you encrypt it client-side, and then it will put it on your iCloud or your Google Drive in an encrypted blob.

Steve: Right, as an encrypted blob, right.

Leo: Which you then can download and so you can move it around. Both you and I have this problem, well, me maybe more so than you, of moving from device to device like a butterfly sampling nectar. And so I have to do this all the time. So I for a long time...

Steve: That would be a good description, yes, Leo.

Leo: Yes. A long time I used Authy, but Authy has that - it does encrypt, but has a disadvantage of storing it in the cloud, their cloud. So I like this 2FAS. And it is open source, which I like, so it's free.

Steve: Good.

Leo: Anyway.

Steve: Apparently multiplatform?

Leo: Yes. Two good choices, though. And, yeah, some day Google says we're going to add end-to-end. And when they do...

Steve: That will be great.

Leo: Yeah. Bran said on Twitter: "The extra protection offered by end-to-end encryption was set aside to balance against 'the cost of enabling users to get locked out of their own data without recovery.'" Which is always the excuse for not using encryption; right?

Steve: So in my drawer I have all of my QR codes printed.

Leo: You print them out, and you put them in a notebook; right?

Steve: That's right.

Leo: Still doing it? Yeah.

Steve: They're in a safe place. And if it ever comes to the point where I need to set up a new authenticator, not a problem. I just scan the QR codes once again, and we're back in business. So the other thing to look for is an authenticator that will allow you to do that because it is nice to have hard copy backup.

Leo: I agree. I agree.

Steve: I have encrypted my most important accounts with two-factor authentication. And it's the right thing to do.

Leo: Oh, god, yes.

Steve: Okay. So I suppose we should not be surprised that Israeli law enforcement is apparently using their own homegrown NSO Group's spyware to spy on their own citizens, in response to reports in Israeli media, which claim that their police have been using a reduced-strength version of the NSO Group's Pegasus spyware known as Saifan.

Leo: Saifan Lite.

Steve: I guess, yeah, I guess you pay less for it, Leo, if it uses a reduced-strength. You don't need the government or the military-grade Pegasus. So you're just some cops, so we're going to give you the reduced...

Leo: Reduced grades.

Steve: Yeah, Pegasus Lite, Saifan, to target activists, business figures, reporters, and politicians. And so, perhaps to save face, the Israeli government has announced the formation - oh, we'll be so glad to hear this, Leo - of a commission.

Leo: Oh, well.

Steve: To probe into the use of spyware by police forces to hack the smartphones of Israeli citizens. Which made me think, isn't the definition of a "commission" the place where sensitive political issues go to die?

Leo: Yes. So they actually admit that they're hacking politicians' and reporters' smartphones. That's stunning. At least they're using the Lite version.

Steve: Yes. The announcement of the commission makes everyone happy, and no one can say that the government isn't doing anything. "Hey! There's a commission for that." But then, after a few years of inaction, it will be quietly disbanded. In any event, if the Israeli media reports are accurate, it was interesting that Israel's own police are also getting in on the act. But the privacy of activists, business figures, reporters, and politicians is being breached.

Leo: Unbelievable.

Steve: Yeah.

Leo: Oh, my gosh.

Steve: Unfortunately, too believable.

Leo: Imagine if the, I mean, I'm sure our government does it, too. But imagine if they admitted that. The uproar.

Steve: Ooh.

Leo: Yeah, the FBI's monitoring the smartphones of reporters and activists and politicians. Ay ay ay.

Steve: Yeah. So rarely have I wanted a Russian translation of anything. But the article in Russia's Kommersant news was written in Russian, and I didn't see any easy way to translate it into English. And by the way, Leo, those Russian characters are really...

Leo: Cyrillic.

Steve: Glyphs and fonts, it's weird. They're weird-looking.

Leo: Yeah.

Steve: But the news is that the Russian government is working on a law to force retailers to pre-install Russian operating systems on all new PCs sold in the country, instead of Windows. The first wave of feedback claims that this will lead to an increase in laptop and PC prices across Russia. Now, despite efforts to get Russian companies and users to move to Russian operating systems, Windows' market share remained the same in Russia as it was last year. And it would be interesting to see what a Russian operating system looks like. I think it has to be a derivative of Linux. You know, that's the only thing that I can imagine would be feasible in this day and age. We've talked about how you just can't start from scratch and create Russki OS. That just - I don't how you'd do that. But if it's a derivative of Linux, why would it be more expensive than Windows? You know, why would it increase the cost?

Leo: Because capitalism.

Steve: Right. So.

Leo: By the way, you should watch the Apple TV show "Tetris" about the history of Tetris. It's fascinating. Speaking of capitalism.

Steve: Oh, cool. I think we did watch it, Lorrie and I, yeah. It was really fun. I completely agree, it was really neat.

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

So we have a story where everyone did everything right. Everyone acted correctly. A problem was found. It was demonstrated. The details were kept secret, and the underlying flaw was responsibly reported to the product's publisher who, in a somewhat timely manner - and it hadn't really mattered how quickly - fixed the trouble and made an update available to their devices. But despite everyone doing everything exactly right, the bulk of TP-Link routers were almost certainly never updated. And with today's Internet scanner databases, discovering the locations of those routers is no longer difficult.

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

Leo: Don't routers auto-update nowadays? No? Yours does not?

Steve: Mine doesn't. And it's a feature that is making its way. But I'm sure you could, if it's even enabled by default; and of course we know if it's not, it might as well not exist.

Leo: Right.

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

And I know this goes against everything we think and believe about fixing and updating known problems. But if patches cannot reasonably be expected to be applied, then what will happen is what just did. Because, I mean, this is not theoretical. Mirai exploded into all the vulnerable TP-Link routers and pwned them. And they were all enslaved into this Mirai botnet.

And, you know, I appreciate the controversy surrounding this, but I think that generic consumer routers all need to occasionally phone home to check for updates and be willing to take themselves offline for an autonomous update cycle.

Leo: Yes. Yes, I agree. Yup.

Steve: The clear prevalence of bad guys who are now waiting to receive and reverse engineer router patches, coupled with the fact that router owners don't know to patch, I think that tips the balance clearly in favor of all such consumer routers being autonomously self-updating by default. Let the owner turn it off if they want. But ship this thing with that checkmark turned on. And for 99.999% of routers, that's the way it's going to stay. They phone home, and they know what time of day it is. They do it in the middle of the night. And they're also monitoring traffic. So they do it when there's like, no - like they find the sweet spot, the block of time where traffic is minimal, and that's when they go download their firmware. There are all kinds of embedded OS solutions, dead man switches and watchdog timers where, if a firmware upgrade were to fail, the hardware could automatically roll back to the previous firmware.

So, I mean, there are ways to do this safely. We talked about this years ago. There was somebody who was, I think, as I recall, someone you knew, Leo, who was involved in the IoT aspect of this, and looking for a safe way to deal with automatic - with IoT devices being able, being empowered to update themselves. This is a perfect classic example of where we really do need that.

Leo: Yeah. My Ubiquiti system will auto-update. And sometimes it's a pain. You know I have it set to do it at 2:00 in the morning. But we have ceiling-mounted WiFi access points. There's one in the bedroom. The light on it is turned off. But when it's updating, it blinks a bright blue. [Indiscernible] woke me up last month saying "There's something wrong. Something's going on. What's that?" I said, oh, it's just the router updating.

Steve: That's very cool.

Leo: But, yeah, and that's nice. And now one of the problems with Ubiquiti stuff is some of the beta versions are notoriously awful. So I have it set for only stable releases.

Steve: Yes, yes.

Leo: But it does it automatically. And I think that's how - frankly, Stacey says, and I agree with her, don't buy home automation stuff that doesn't auto-update.

Steve: That's really good.

Leo: In this day and age you need it.

Steve: Yup. I'm glad that that's beginning to be the word that is spread. I would not have a consumer router on the public Internet. All of mine are behind a pfSense firewall. So I have a separate box that is in front. And actually I need that because I need - I do a bunch of crazy stuff with port mapping in order to link my two sites and get around Cox's blocking of ports that are useful to have open.

Leo: Right. And does pfSense auto-update? It must; right?

Steve: No, it doesn't. It will check for updates, but it won't do it by itself.

Leo: Yeah, I guess that's the theory is we need to be 100% uptime.

Steve: Yeah, exactly. And I think that's something we're going to have to get over. And then of course the other problem is what if the update fails, and the router stumbles, and now you're offline, and you have no connectivity.

Leo: Right.

Steve: Well, again, there are ways to roll back from a failed update.

Leo: Yeah. And Ubiquiti will do that automatically.

Steve: Perfect, perfect.

Leo: I'm very happy with this gear. It's been very good.

Steve: So we have a bit of happy news from Intel. We first need to know what Intel's TDX is. TDX stands for "Trust Domain Extensions." Intel describes it as: "Intel Trust Domain Extensions is introducing new architectural elements to deploy hardware-isolated virtual machines called Trust Domains. Intel TDX is designed to isolate VMs from the virtual-machine manager/hypervisor and any other non-TD software on the platform to protect Trust Domains from a broad range of software." So anyway, it's further virtual walls in order to control security.

They said: "VM isolation with Intel TDX is a key component of Intel's Confidential Computing portfolio, which also includes application isolation with Intel SGX and trust verification with our upcoming service code-named Project Amber. Confidential

Computing uses hardware to protect data in-use from a wide variety of threats, and enables organizations to activate sensitive or regulated data that may have otherwise been locked down and idle." So, okay. So it's more security stuff that they're building into their baseline hardware.

After this announcement, Intel brags that before releasing this new tech to the world, they ran it through a very useful security gauntlet. They wrote three points. First: "In our first-ever pre-release activity, we also took Intel TDX through Project Circuit Breaker, part of Intel's Bug Bounty, where we challenged a community of elite hackers to find bugs in some of our top technologies. Using simulation software, the community went through two rounds of bug hunting over several months, earning bounties to help us find potential vulnerabilities so we could mitigate them.

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

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

Okay. So the good news is that even though the authors of Intel's code were as sure as any code authors ever are that their code was correct, they nevertheless subjected it to pre-release third-party scrutiny. Naturally, Intel put a happy face on the results, saying that it had succeeded in improving the code quality. What we learn from other sources is that, indeed, vulnerabilities were uncovered during the security audits that could have resulted in arbitrary code execution, cryptographic weaknesses, and denials of service. So I hope that somebody who's pulling the strings over there recognized fully how much benefit they got from this, and that this becomes standard practice because it's a great idea.

And in what Intel wrote, they didn't say something that I did see elsewhere which is they provided the source code. They made the source code available to Google's Project Zero guys so they could really take a look at it and weren't being forced to reverse engineer it and just guess at what was going on.

On the flipside, an academic paper was just published titled "Timing the Transient Execution: A New Side-Channel Attack on Intel CPUs." And since we've entered the world of "yet another side-channel information leakage from Intel CPUs," I'm not going to spend undue time digging into this one. But in case it might wide up being important, which hopefully is unlikely, I wanted to at least share the brief description of these authors. They were a bunch of Chinese researchers, but located domestically. But you'll hear in their description that some of their word choices are a little confusing. But still, we can see what's going on.

They wrote: "The transient execution attack is a type of attack leveraging the vulnerability of modern CPU optimization technologies. New attacks surface rapidly. The side-channel is a key part of transient execution attacks to leak data. In this work, we discover a vulnerability that the change of the EFLAGS register in transient execution may have a side effect on the Jcc" - that's Jump on condition code - "instruction after it in Intel CPUs." And EFLAGS is just, you know, all processors have a status register, a FLAGS register. "E" stands for Extended because it used to be 16 bits long, and now it's 32 bits long.

Anyway, they said: "Based on our discovery, we propose a new side-channel attack that leverages the timing of both transient execution and Jcc instructions to deliver data. This attack encodes secret data to the change of register which makes the execution time of context slightly slower, which can be measured by the attacker to decode data. This attack does not rely on the cache system and doesn't need to reset the EFLAGS register manually to its initial state before the attack, which may make it more difficult to detect or mitigate.

We implemented this side-channel on machines with Intel Core i7-6700, i7-7700, and i9-10980XE CPUs. In the first two processors, we combined it as the side-channel of the Meltdown attack, which could achieve 100% success leaking rate. We evaluate and discuss potential defenses against the attack. Our contributions include discovering security vulnerabilities in the implementation of Jcc instructions and EFLAGS register, and proposing a new side-channel attack that does not rely on the cache system."

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

And the reason not to get too overworked about this is that it does require that hostile code already be present. Certainly in our own personal workstations, if something's in them that's evil, it's already too late. The reason concern is in all of the virtualization which is going on now where you might have multiple companies' systems sharing a common set of hardware. And there it's cross-OS or cross-VM leakage that you need to be concerned about.

Leo: There have never been any in the wild, that we know of, examples of these side-channel exploits, like Meltdown and Spectre.

Steve: No. Still...

Leo: It's hard to do.

Steve: Yes, I mean, even Heartbleed, you know, it required that you pound on the server for a long time, and maybe you got something. And remember when Heartbleed first came out, the discoverers recognized and called it a "theoretical problem," but doubted it could ever actually happen. And then it did.

Leo: Right.

Steve: So, yeah. So the real problem here is that what we're seeing Intel being forced to do is they're having to back out of very useful performance optimizations driven by the world that doesn't want to have any known possibility of problems.

Leo: Well, and the problem is mostly on shared servers; right? You said you had...

Steve: Only. Only on shared servers.

Leo: Only on shared servers.

Steve: Yeah.

Leo: So if you're, you know, it's not your home machine you have to worry about.

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

Leo: cURL is the single probably most useful program for interacting with the Internet there is. I mean, it's incredible.

Steve: Yes, yes. So here's what Daniel posted to Mastodon. His first post said: "Do NOT." NOT in all caps. "Do NOT. I repeat. Do NOT remove curl.exe from your Windows System32 folder to silence a (stupid) security scanner. It will lead to tears and sorrows. And if you do, please don't ask me for help when you've broken your Windows install. I can't fix that."

And he followed up the tweet with: "Why do people remove it? Because NVD has exaggerated a cURL security flaw to an inflated level, and now 'security scanners' [he has in quotes] insist that the bundled cURL executable has a 'high severity' [again in quotes] security flaw and scaremongers people into removing it."

Leo: Wow.

Steve: "And then they realize Windows Update refuses to work."

Leo: Oh, interesting.

Steve: And then he finishes: "Are we sure this is the best we can do?"

Leo: So Windows Update uses cURL, I guess.

Steve: Apparently some facet of Windows Update does, yeah. And I didn't realize it's sitting there, a curl.exe in the System32 folder.

Leo: Wow.

Steve: Yeah. Very cool.

Leo: Well, certainly it's on all Linux distros and all Macs.

Steve: Oh, yeah.

Leo: Yeah.

Steve: Yeah, I'm a big Wget user. So that's a...

Leo: I always replace it with Wget because it's easier, but...

Steve: Yes. And just a little tip for Wget fans. There was something I needed to download. Actually I'm sure it was some network driver, long forgotten, for a card that a motherboard had that I needed to get on the network in order to do network-based debugging of SpinRite. So wherever it was, I think it was on IBM, actually, on IBM.com there was an archive. When I clicked on it, no luck. It was FTP. And then I remembered.

Leo: Yes.

Steve: That Wget does FTP.

Leo: Yeah, yeah.

Steve: So.

Leo: cURL supports every possible - you can do posts and gets and everything with cURL. It's awesome.

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

So I think this is potentially very cool and useful. However, increasingly real-world evidence is suggesting that AI, at least in the form of our well-known ChatGPT, always sounds absolutely convincing and authoritative while being completely wrong with its facts. So at least at this early stage, anything and everything that's produced should be regarded with skepticism and carefully vetted and verified. By all means, you know, see what the AI has to say, but then go do your homework. You know, use that as a starting off point and verify it yourself.

On the other hand, since this is a Google effort, and they know their technology, it would be truly cool to have something where you could update this, and it would provide an automated analysis of whatever malware you provided. So I hope that this gets better and continues to grow in the future. Sounds like a good idea.

Leo: Yay.

Steve: And Leo?

Leo: Well, well, well. Here we are.

Steve: That's our show.

Leo: Once again at the very end. Sadness rules the land.

Steve: And remember, I think it was on Windows 95 there was that really bizarre comment that would come up. Maybe it was the paper clip? And it said: "Never run with scissors?"

Leo: No, I don't remember that.

Steve: Oh, yeah, it was like...

Leo: That sounds like Clippy, Clippy always looking out for you, yeah.

Steve: I just wanted to tell anyone, never be in a hurry when you're shaving with a razor.

Leo: Oh, my gosh. It's like Jack the Ripper hit you.

Steve: Because, yeah, well, it looks worse than it is, you know, it's just a paper cut. Just a flesh wound. I'll be fine.

Leo: You used to be an adventurer. Then you took a razor to the chin, and it's all over.

Steve: I haven't cut myself a lot for a long time.

Leo: Steve, always a thrill, always a pleasure. It's great to see you. I'm glad you've got your power back at your headquarters, your world headquarters.

Steve: Yes. Power is at both locations, so I will be retiring this emergency - we call this, what was it on the Enterprise? It was the Auxiliary Bridge.

Leo: Auxiliary Bridge shutting down for the week.