



Hiding School Cyberattacks

Description: New "SparkCat" secret-stealing AI image scanner discovered in App and Play stores. The UK demands that Apple do the impossible, decrypt ADP cloud data. France moves forward on legislation to require backdoors to encryption. Firefox moves to 135 with a bunch of useful new features. The Five Eyes alliance publishes edge-device security guidance. Six Netgear routers contain CVSS 9.6 and 9.8 vulnerabilities. Sysinternals utilities allow malicious Windows DLL injection. Google removes restrictive do-gooder language from AI application policies. "AI Fuzzing" successfully jailbreaks the most powerful ChatGPT o3 model. Examining the well and deliberately hidden truth behind ransomware cyberattacks on U.S. K-12 schools.

High quality (64 kbps) mp3 audio file URL: <http://media.GRC.com/sn/SN-1012.mp3>

Quarter size (16 kbps) mp3 audio file URL: <http://media.GRC.com/sn/sn-1012-lq.mp3>

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about that malware that crept into the Apple App Store just a couple of weeks ago. The UK says Apple has to put a backdoor in its encryption. Steve's opinion on that. And we'll talk about how common it is for schools in the United States to hide the fact that they've been ransomware'd, and why they do it. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 1012, recorded Tuesday, February 11th, 2025: Hiding School Cyberattacks.

It's time for Security Now!, the show where we cover the latest security news, privacy news, encryption news, and sprinkle in a little bit of stuff about science fiction and TV shows and whatever else this guy right here is into at the moment. Steve Gibson, our polymathic master of ceremonies. Hi, Steve.

Steve Gibson: Hello, my friend. It's great to be with you again for, now, we're out of the binary episodes. We had one zero, one zero. Well, we had of course 1000. Then we had 1001. Then we had 1010, 1011. Now we're, unfortunately, until we get to 1100, which is not that far...

Leo: No, it's only 88 episodes away. We could do it.

Steve: I think we're going to still be here.

Leo: That's less than two years. We will definitely be here for 1100. Until then, back to decimal.

Steve: So I ran across, while I was just catching up on news, and we've got a bunch of interesting news, ran across an entity I'd never heard of called The 74. The 74 is as in 74 million, which is the number of kids in K-12 lower education schools. And so this is an organization that basically represents their interests. And it's nonpartisan. It's straight down the middle. It's, I mean, and they have a code of ethics that their reporters follow.

A reporter did some amazing investigative journalism revealing the degree of the problem that this country has with hidden educational school cyberattacks, which are going on at a much higher rate than is believed because it turns out there are actual coaches that are quite busy coaching educators about how not to reveal the fact that their school has been compromised. And of course this has consequences downstream for the kids whose personal data and disciplinary records, family problems, emotional health, all kinds of personal information is out, and the educators are denying it.

Anyway, I know that a lot of our listeners are parents of school-age kids. I've had feedback from them through the years about this. And there was a lot to say about this. So we're going to spend half of the podcast going through this detailed piece of investigative journalism because, you know, I thought, well, can I summarize this somehow? But what I came away with after reading the entire thing was a much, I mean, a truly deep understanding of the dynamics of this. And unfortunately I'll be concluding, drawing some conclusions which I think everyone will be able to follow by the time we get through this. But so that's our main topic for the day. And if it's something that's not of interest to anybody, well, fine, you know, stop listening after you get to the first hour and a half.

Leo: After you hear all the ads, Steve. After you hear all the ads.

Steve: Oh, that's right. Right. After the last ad we'll let you know you can leave now. You can leave the classroom, children.

Leo: You can leave anytime.

Steve: Students of cybersecurity. But we're going to talk about first SparkCat, which I know you guys talked about over on MacBreak Weekly. SparkCat is the name of the secret-stealing AI image scanner.

Leo: Oh, yes.

Steve: Which has been discovered in the App and the Play stores. Also, I also saw you mentioning the UK's new demands about Apple doing the impossible; and we're going to touch on, you know, the whole advanced data protection issue. We talked about it when Apple first announced it. UK's back again. France is also moving forward on legislation to require backdoors into encryption. And as I've been saying now for a couple years, this is the great question; right? This is why we're glad we went past 999, because this, like, how do you solve this? And I'm going to reiterate the solution that I think exists, which also defuses the arguments about why government needs this stuff.

Firefox has moved to their number 135, and it has a bunch of useful new features, which I don't see all the time. But if you hold, I think it's CTRL+SHIFT+X, up pops ChatGPT. So, and, yeah, they've added chatting to an optional sidebar in Firefox, and a bunch of other stuff we'll talk about. Also the Five Eyes alliance has published their guidance for edge-device security, which I know our listeners who are involved in enterprise environments will like because it's both a checklist and a CYA for them.

Six Netgear routers contain, three each in three sets, CVSS 9.6 and 9.8 vulnerabilities, unauthenticated remote code execution. You want to make sure you don't have any of these six routers. And you're only vulnerable if you also didn't follow the guidance, and you've got remote admin enabled, which of course I'm sure none of our listeners have.

Leo: Oh, never, no.

Steve: But if you did, oh, boy. And one of our favorite classes of utilities, those from Sysinternals, it turns out that most of them allow malicious Windows DLL injection. And apparently Microsoft doesn't care. We'll look at that more closely. Google has removed restrictive do-gooder language from their AI policies, which has raised some questions. And there's an open source, posted up on GitHub, AI Fuzzer, which has successfully jailbroken the most powerful and supposedly guardrail-equipped ChatGPT o3 model. We're going to look at all that. And then we're going to end by examining the well and deliberately hidden truth behind ransomware cyberattacks on the U.S.'s K-12 schools.

Leo: Wow.

Steve: And of course we have a picture for the ages, Leo. This is one where it's like, okay. I think I mentioned it, actually, either last week or the week before, the nominee for the 2025 Darwin Awards.

Leo: You know, I love the Darwin Awards. I haven't thought about them in a long time. Have we got an image for you.

Steve: Well, I think we remember the one where someone, like, strapped a skateboard to a jet engine or something, I mean, there was, like, what? Anyway...

Leo: It's called the Darwin Awards because most of the people who do these experiments do not survive. And it is survival of the fittest, or at least the least risk-taking. Good. I can't wait. I haven't looked at it yet. We'll look at it together in just a moment.

Steve: Well, it takes a minute to sort of parse the picture. And then you think, OMG.

Leo: Okay. Oh, I can't wait. I love the puzzle ones. That's good. All right. Time for our Picture of the Week. Now, I want to scroll up. And then after I do I will look at it for a minute, and then I'll let you all - nominee for the Darwin Awards. Oh, my god. That's got to be staged. If it's not, then let's do a little requiem for this guy.

Steve: Oh, yeah. So...

Leo: I understand the need to party. I certainly do.

Steve: So we start with a very large inflatable backyard swimming pool, looks like it could hold about 50 people. So, you know, big blue perimeter filled with water. Now, apparently these guys didn't want to get out of the pool in order to, I don't know, enjoy some grilled meat of some sort.

Leo: It's a party in the pool. They've got the beers. They've actually taken a table and put it in the pool.

Steve: They put a table, yes. Yes. Now, the problem is that, rather than using briquettes of any sort, they thought, well, let's just, you know, we don't have a wood-burning grill.

Leo: We have an electric griddle.

Steve: We've got an electric grill.

Leo: Yeah.

Steve: Now, we need to get power to this grill, but the cord's not long enough to reach from the grill to past the perimeter of this round backyard pool. So we need to use an outlet strip.

Leo: Sure, as one does.

Steve: Which I order to, like, you know, bridge from the grill to outside of the pool.

Leo: But these guys are very safety conscious. They know you can't submerge that strip.

Steve: Oh, no, no, that would not be good, no.

Leo: No.

Steve: So they took some floating flip-flops, some floating...

Leo: Those are shower slippers.

Steve: Yes, some sandals. And they stuck them on each end of the power strip to float them in the middle of the pool. Meanwhile, we see the cord from the power strip going over to the edge, where it meets some sort of a, looks like a wood block like a door wedge.

Leo: I don't know what they're doing with that.

Steve: Connected to another extension cord.

Leo: Maybe that's so that it will float if it falls in. This is ridiculous.

Steve: Oh, boy. Yes.

Leo: So what, you know, just out of curiosity, what would happen if that power strip fell in the pool?

Steve: Okay. Now, to the credit, giving credit to our listeners, I finished putting the show notes together yesterday, late afternoon. So I did the mailing. It went out to 16,100 of our listeners who have subscribed to the weekly mailing. So I've had feedback about this picture since then. Many of our listeners said, well, you know, they're in a plastic rubber pool.

Leo: It's insulated.

Steve: They're not connected to ground.

Leo: They're not grounded.

Steve: That's right. That's very much like...

Leo: Plus it's a surge strip. It probably has a fuse.

Steve: Well, it's very much like - one listener drew the analogy to the bird that lands on the high-tension wire. It lands on the wire. It doesn't turn into barbecued bird immediately. You know? It doesn't know what's going on. Maybe its teeth hurt a little bit. I'm not sure. But basically it's going to survive. So the problem would be, assuming that this weird sandal power strip thing capsizes in the pool, it would just blow the fuse. I mean, back at the house a fuse would blow. The coffee pot would turn off in the kitchen.

Leo: Or the surge protector would blow, yeah. But you know what could be trouble is if, at the same time as it sank, somebody stepped out of the pool and had one leg in the pool.

Steve: Oh.

Leo: And one leg on the ground.

Steve: Then, yes, then you're the bird that lands, straddles two wires, and we remember what happened with Wile E. Coyote on "The Road Runner."

Leo: Yes.

Steve: It turns into a little poof, just a little shadow of its former self, little crispy sticks that then drops down to the bottom of the ground.

Leo: I want to see more Darwin Award pictures.

Steve: So what you need to do basically...

Leo: Burke says, by the way, and Burke has probably been electrocuted many times in the studio, he says the wire is the ground. So maybe, I mean, it could just go right back to the ground through the wire; right?

Steve: Well, now, this guy who's further away, who seems to think this is quite entertaining...

Leo: Hysterical? Yeah.

Steve: He's smiling. It looks like he's pushing the edge of the pool down because he's...

Leo: Also a mistake.

Steve: Yeah. And so if any water is running out past him, that's running to ground, and he could be helping to close the circuit.

Leo: Holy cow.

Steve: Yeah.

Leo: Okay.

Steve: People have survived being struck by lightning, apparently. Again, their teeth hurt.

Leo: Do not throw a toaster in the pool. I'm just saying.

Steve: No.

Leo: Or a grill.

Steve: Don't do that. Actually, don't do any of this, kids.

Leo: No.

Steve: Bad idea. Thus the Darwin Award, although it looks like these people probably already have reproduced, so this won't be limiting their future ability to propagate this foolishness.

Leo: Oh, dear. Oh, dear.

Steve: So, as we know, the United States has shunned the Russian cybersecurity firm Kaspersky over understandable, if unfair, concerns of the potential for Russian influence, which would be truly devastating if Kaspersky were to ever turn malicious, given how much of their software was phoning home from inside the U.S. Again, it's unfortunate that this happened. Kaspersky is nevertheless continuing to contribute their important security research to the world, and their publication last Friday about their discovery of a new trojan, which they've dubbed "SparkCat" (S-P-A-R-K-C-A-T) is another example of Kaspersky's continuing value to everyone, even though we don't want to trust them anymore. Which, as I said, is too bad.

So I want to share the details of their discovery because it should put everyone on notice of the way malware is evolving. And I'm sure that the fact that it illustrates the potential for the abuse of Microsoft's own screen-scraping "Recall" technology will not be lost on any of our listeners who dislike the idea of having their PCs constantly being scraped and archived because that's sort of the way this thing works.

Kaspersky's piece is titled "SparkCat trojan stealer infiltrates App Store and Google Play, steals data from photos." And they follow that with the tag: "We've discovered apps in the official Apple and Google stores that steal cryptocurrency wallet data by analyzing photos." So here's what they published last Friday. And, you know, it strikes me this is sort of the evolution of the earlier cryptocurrency stealers which were monitoring people's clipboards; right? They were, like, polling the Windows clipboard because it would be like, it would be natural for a Windows user who is wanting to send some cryptocurrency somewhere, you know, it's like, you know, you're paying somebody through bitcoin, and they say, "Here's our address. Send us X bitcoin." And so, you know, the addresses are crazy.

So you copy that address with your mouse, you know, hit CTRL+C to copy it. Then you go over to your bitcoin wallet where you want to send bitcoin to an address, and you paste it. Well, you don't even - at no point do you try to, like, study that gibberish of a bitcoin address. You just blindly copy and paste. Of course, as we know, cryptocurrency stealers, first-generation ones, they would watch for the arrival of a bitcoin address and quickly substitute their own so that the address you pasted was not the address you had copied. And you ended up sending them the money that you were intending to send

somewhere else. And then, you know, after a while you contact the original group and say, "Hey, where is my thing that I ordered?" And they go, "Where is the money that you're supposed to send?" And you say, "I sent you the money." It was like, "Well, no." Anyway, bad guys got it.

So here's the evolution of that. Kaspersky said: "Your smartphone gallery may contain photos and screenshots of important information you keep there for safety or convenience, such as documents, bank agreements, or seed phrases for recovering cryptocurrency wallets. All of this data can be stolen by a malicious app such as the SparkCat stealer we've discovered. This malware is currently configured to steal crypto wallet data, but it could be easily repurposed to steal any other valuable information." And again, it's like, it's one thing to be careful about not putting, like, taking a photo of something with your phone, despite the fact that we're told how secure that is. But it is really creepy if your Windows desktop was doing that continuously.

They said: "The worst part is that this malware has made its way into official app stores, with almost 250,000 downloads of infected apps from Google Play. Although malicious apps have been found in Google Play before, this marks the first time a stealer trojan has been detected in the App Store. How does this threat work, and what can you do to protect yourself?

"Apps containing SparkCat's malicious components fall into two categories. Some, such as numerous similar messenger apps claiming AI functionality, all from the same developer, were clearly designed as bait. Some others are legitimate apps: food delivery services, news readers, crypto wallet utilities." They said: "We don't yet know how the trojan functionality got into these apps. It may have been the result of a supply chain attack, you know, where they broke into the apps' developers and injected their library," they said, "where a third-party component used in the app was infected. Alternatively, the developers may have deliberately embedded the trojan into their apps.

"The stealer analyzes photos in the smartphone's gallery; and to that end, all infected apps request permission to access the user's photos. In many cases, this request seems completely legitimate. For example, the food delivery app ComeCome requested access for a customer support chat right upon opening this chat, which looked completely natural. Other applications request gallery access when launching their core functionality, which still seems harmless. After all, you do want to be able to share photos in a messenger; right?

"However, as soon as the user grants access to specific photos or the entire gallery, the malware starts going through all the photos it can reach, searching for anything it might find valuable. To find crypto wallet data among photos of cats and sunsets, the trojan has a built-in optical character recognition module based on the Google Machine Language Kit" - the Google ML Kit "a universal machine-learning library.

"Depending upon on the device's language settings, SparkCat downloads models trained to detect the relevant script in photos, whether Latin, Korean, Chinese, or Japanese." So multilingual. "After recognizing the text in an image, the trojan checks it against a set of rules loaded from its command-and-control server." Thus its function, what it's doing when it finds things, can be varied on the fly. They said: "In addition to keywords from the list - for example, 'mnemonic' - the filter can be triggered by specific patterns such as meaningless letter combinations in backup codes or certain word sequences in seed phrases.

"The trojan uploads all photos containing potentially valuable text to the attackers' servers, along with detailed information about the recognized text and the device the image was stolen from." So it's serving sort of as a frontend filter, doesn't want to swamp these nefarious creators, the developers, with everything, you know, every photo

in everyone's phone who downloads it. So it does upfront filtering to determine if anything is interesting, no sunsets and cat pictures.

They said: "We identified 10 malicious apps in Google Play, and 11 in the App Store. After notifying the relevant companies, and before publishing this, all malicious apps had been removed from the stores. The total number of downloads from Google Play alone exceeded 242,000 at the time of analysis, and our telemetry data suggests that the same malware was available from other sites and unofficial app stores, as well.

"Judging by SparkCat's dictionaries, it is trained to steal data from users in many European and Asian countries, and evidence indicates that attacks have been ongoing since at least March of 2024." So this is coming up on a year old. "The authors of this malware are likely fluent in Chinese. More details on this, as well as the technical aspects of SparkCat, can be found in the full report on SecList." Which is the Secure List is where Kaspersky posts all of their technical stuff.

So under "How to protect yourself from OCR trojans," they write: "Unfortunately, the age-old advice of only download highly-rated apps from official app stores is a silver bullet no longer. Even Apple's App Store has now been infiltrated by a true infostealer, and similar incidents have occurred repeatedly in Google Play. Therefore, we need to strengthen the criteria here. Only download highly rated apps with thousands, or better still, millions of downloads, published at least several months ago. Also, verify app links in official sources (such as the developers' website) to ensure they're not fake, and read the reviews, especially the negative ones."

They said: "You should also be extremely cautious about granting permissions to new apps. Previously, this was primarily a concern for 'Accessibility' settings, but now we see that even granting gallery access can lead to the theft of personal data. If you're not completely sure about an app's legitimacy (for example, it's not an official messenger, but a modified, like enhanced version), don't grant it full access to all your photos and videos. Grant access only to specific photos, and only when necessary. Storing documents, passwords, banking data, or photos of seed phrases in your smartphone's gallery is highly unsafe."

Leo: Yeah, don't do that.

Steve: You start with not doing that, yes.

Leo: Get a password manager.

Steve: Yup. They said: "Besides stealers such as SparkCat, there's also always the risk that someone peeks at the photos, or you accidentally upload them to a messenger or file-sharing service. Such information should be stored in a dedicated application." To your point, Leo, exactly that.

And they said: "Finally, if you've already installed an infected application (the list of them is available at the end of the SecList post), delete it and don't use it until the developer releases a fixed version. Meanwhile, carefully review your photo gallery to assess what data the cybercriminals may have obtained. Change any passwords and block any cards saved in the gallery. Although the version of SparkCat we discovered hunts for seed phrases specifically, it's possible that the trojan could be reconfigured to steal other information. As for crypto-wallet seed phrases, once created, they cannot be changed.

Create a new crypto wallet, transfer all your funds from there, and completely abandon the compromised one."

Leo: Wow. We should say that Apple deleted, has presumably killed - well, I don't know.

Steve: If they've retroactively killed the...

Leo: Yeah, they killed the apps in the store, but maybe they didn't use the kill switch. They have a kill switch to delete apps unsafe like that. Maybe they...

Steve: And Kaspersky's point is that, even if they did, it's worth, I mean - and the good news is these are not mainstream apps. You know, ComeCome, which is some Chinese food delivery service...

Leo: Oh, okay.

Steve: It's, you know, it's not something that a lot of people are probably going to have. On the other hand, 242,000 people had downloaded apps that had this in it from Google Play. And so Kaspersky's point is it's worth auditing the photos that you have to see what they may have gotten and get proactive. Because if you've got a bunch of bitcoin, and you're storing your recovery phrases in a photo in your photo library, first of all, bad idea. But secondly, you know, if you were to find that in your photo library, good idea to just create a new wallet and move everything over there and just don't do that again. So anyway...

Leo: Now, if somebody stole the password or the recovery phrase from my wallet, I would really appreciate it if you'd just let me know because I'll give you 10%.

Steve: You would call it a commission if they - that's right.

Leo: Exactly.

Steve: Okay. So I linked to Kaspersky's full technical report for anyone who wants to dig into this more deeply. I'll go one step further than Kaspersky has in my advice. Just as is true with today's web browsers whose users have demanded openness in the form of browser add-ons, the same openness has been demanded and received from mobile phone manufacturers. Unfortunately, there are bad guys in the world who profit from victimizing others. The other thing we've seen is that despite the best efforts of those managing the add-ons that are available for our browsers and our phones, malicious applications still manage to sneak in. The good news is that one thing we've seen over and over is that the least secure and malice-prone applications are typically, you know, again, as I've called them, I guess I would call them sort of gratuitous editions. They're apps that everyone can live without.

So their victims tend to be people who download anything that comes along that looks even remotely interesting. I mean, they've completely lost control of their phones. They

don't know where any of their apps are. They just scroll endlessly trying to find an icon for something that they're looking for. You know, they have no appreciation for the fact that there is a non-zero chance that the creator of any given app may have malicious intent or may not, but used a malicious library without knowing it. The point is non-zero, which tells us, just the law of statistics and probability and numbers, the more apps you have, where each one has a non-zero chance of being a problem, the greater the total problem. It only takes one in order to create a leak.

So my advice is always to keep this in mind when deciding whether you really need the app you're considering. And because our devices' manufacturers have done everything they can to give us the tools to restrain what apps can do, even after they're resident in our devices, be parsimonious with the access permissions that apps are granted. And I know this is tricky since apps will be cleverly designed to need the permissions they wish to abuse, but at least always question their need.

And Leo, you know, it just is a fact that we're seeing arguably more of this today than when this podcast began 20 years ago.

Leo: Oh, yeah. Absolutely. Because there's more money to be had; right?

Steve: Yeah.

Leo: Everything's on our phones nowadays.

Steve: And more devices available. I mean, everybody has one. You don't see - I see people, you know, everybody sitting in a restaurant is staring at their individual phones. I don't know how people don't fall off the curb. They're walking down the sidewalk.

Leo: I know, it's really bad.

Steve: Staring at their phones. I think, what is it? What are you doing?

Leo: It's amazing. Well, yeah. They're desperately avoiding any boredom or feelings or knowledge of the world. They're narcotic. They're narcotizing themselves. I do think, and you said something really important, and I really would underscore this, and I always said on the radio show, install the fewest possible apps. You know?

Steve: Yes.

Leo: On your desktop, on your laptop, on your iPad, on your phone, the fewer the apps, the better because every app raises the specter of a security flaw or just a bug.

Steve: Also, we are seeing the built-in apps slowly subsuming the functionality. I used to use a - I had a really cool perspective correction app that I was using. Now...

Leo: Camera does it.

Steve: ...it's built in. We used to - there was never an Edit button in the beginning for photos. Now you push Edit, you can rotate them, you can fix the perspective, do all this kind of stuff. So, you know, you don't need third-party apps to do that.

Leo: You'd probably, on an iPhone or a good Android phone like a Google Pixel, could get away without any apps.

Steve: I think that's really the case.

Leo: And that would be a lot safer, for sure.

Steve: I should mention one of your guests, I think it might have been Alex, mentioned Foreca, F-O-R-E-C-A.

Leo: Yeah, we've been talking about that for a long time as the best weather app.

Steve: I absolutely, I just, I've been wanting to mention to our listeners. Is it multiplatform?

Leo: Yeah.

Steve: Is it available on Android?

Leo: Yeah. It's everywhere, yeah.

Steve: It is so good. They asked me after a year, I think, they said: "Can we have a little more money?" I said yes. Because, you know, I want them to keep it the way it is. Anyway, it's short for Forecast, F-O-R-E-C-A. It made me think of it because I'm not - I don't think Apple's weather thing, you know...

Leo: It's pretty barebones.

Steve: ...holds a candle to Foreca.

Leo: It's pretty barebones.

Steve: So there's an example of something where you really - okay, yeah. But you can trust these guys.

Leo: Yeah, I think you can trust them. A lot of weather apps actually use Foreca as the back end. I think my CARROT Weather is using - or at least you can choose a forecast backend. But their app is quite good.

Steve: And you can look at satellite and radar.

Leo: Yeah.

Steve: And it does take some getting used to. It is very information dense. And it's also customizable, what things you do care about and you don't. I don't care about wind that much, so it takes up space on my screen. Get rid of wind, but I do want rainfall. And, boy, like to know what time of day something's going to happen. Anyway, just an unsolicited note that I've been meaning to mention it because I just keep really liking it.

Leo: They're on the web. They're on the Google Play Store, the Apple Store.

Steve: The web. Ahhh.

Leo: Yeah, you can use the web version.

Steve: Big screen.

Leo: Yeah.

Steve: Nice.

Leo: And it has the videos and everything, too.

Steve: Yeah.

Leo: Which is really pretty cool. I like it. Look at that. All you need is a green screen, Steve, and you can do your own weather report.

Steve: Yes. I don't - okay. I was going to say something off-color, but...

Leo: No, no.

Steve: Not a weather girl.

Leo: That's good.

Steve: And on that note, let's take a break, and then we're going to come back and talk about the UK's news demand for Apple's encrypted data.

Leo: Oh, I do really want to hear what you have to say about that. We knew it was coming. It just - it's finally happened; right?

Steve: Yes, well, it's like they keep trying; right? They just, like, they keep hitting this immovable wall. I have an idea, though.

Leo: Oh. Now I'm liking it. Stay tuned. Steve has an idea. All right, Steve. I'm very curious. They call it the Snoopers' Charter, you know.

Steve: Yeah.

Leo: The Investigatory Powers Act.

Steve: So last Friday the news broke that the United Kingdom was demanding that Apple provide access to its users' cloud data. I received links from our listeners to stories of this in The Register, The Guardian, and the BBC. These reports were all picking up the news which was first reported in The Washington Post. And The Post provided the best coverage of all. So let's turn to the source for the whole story. So here's what we know from The Washington Post's reporting.

They said: "Security officials in the United Kingdom have demanded that Apple create a backdoor allowing them to retrieve all the content any Apple user worldwide has uploaded to the cloud, people familiar with the matter told The Washington Post." Okay. So again, all the content any Apple user worldwide has uploaded to the cloud. Good luck with that. But this is what they say they want.

"The British government's," writes The Post, "undisclosed order, issued last month, requires blanket capability to view fully encrypted material, not merely assistance in cracking a specific account, and has no known precedent in major democracies. Its application would mark a significant defeat for tech companies in their decades-long battle to avoid being wielded as government tools against their users, the people said, speaking under the condition of anonymity to discuss legally and politically sensitive issues.

"Rather than break the security promises it made to its users everywhere, Apple is likely to stop offering encrypted storage in the UK, the people said. Yet that concession would not fulfill the UK's demand for backdoor access to the service in other countries, including the U.S. The office of the Home Secretary has served Apple with a document called a 'technical capability notice,' ordering it to provide access under the sweeping UK" - I always trip up on that - "UK Investigatory Powers Act of 2016...

Leo: Nice, well done.

Steve: ...which authorized law enforcement to compel assistance from companies when needed to collect evidence, the people said.

"The law, known by critics as," as you said, Leo, "the Snoopers' Charter, makes it a criminal offense to reveal that the government has even made such a demand. An Apple spokesman declined to comment." After all, they can't reveal that. "Apple can appeal the UK capability notice to a secret technical panel, which would consider arguments about the expense of the requirement, and to a judge who would weigh whether the request was in proportion to the government's needs. But the law does not permit Apple to deny complying during an appeal." Meaning you can't use the appeal to delay the order, which of course means that the information that the UK would want would already be in their possession. Even if Apple were to win the appeal it'd be too late. So this is a mess.

"In March," writes The Post, "when the company was on notice that such a requirement might be coming," so almost a year ago, "it told Parliament: 'There is no reason why the UK government should have the authority to decide for citizens of the world whether they can avail themselves of the proven security benefits that flow from end-to-end encryption.'" "The Home Office said Thursday that its policy was not to discuss any technical demands. Their spokesman said: 'We do not comment on operational matters, including, for example, confirming or denying the existence of any such notices.'" In other words, no comment.

"Senior national security officials in the Biden administration had been tracking the matter since the UK first told the company [Apple] it might demand access, and Apple said it would refuse. It could not be determined whether they raised objections to Britain. Trump White House and intelligence officials also declined comment.

"One of the people briefed on the situation, a consultant advising the United States on encryption matters, said Apple would be barred from warning its users that its most advanced encryption no longer provided full security. The person deemed it shocking that the UK government was demanding Apple's help to spy on non-British users without their governments' knowledge."

Leo: This is really important. It includes us.

Steve: Yes. Yes. And a former White House security adviser confirmed the existence of the British order. So in the reporting, The Washington Post did their due diligence, and they got multisource confirmation that this is all happening and has happened. "At issue," they finish, "is cloud storage that only the user, not Apple, can unlock. Apple started rolling out the option, which it calls Advanced Data Protection, in 2022. It had sought to offer it several years earlier, but backed off after objections from the FBI during the first term of President Donald Trump, who pilloried the company for not aiding in the arrest of 'killers, drug dealers and other violent criminal elements.'" The service is an available security option" - we're talking about Advanced Data Protection - "for Apple users in the United States and elsewhere.

"While most iPhone and Mac computer users do not go through the steps to enable it" - because it's not enabled by default - "the service offers enhanced protection from hacking and shuts down a routine method law enforcement uses to access photos, messages, and other material. iCloud storage and backups are favored targets for U.S. search warrants, which can be served on Apple without the user knowing."

So, and just for the record, remember it's often not a question or choice about whether you want ADP enabled. I'd love to have it enabled, but I cannot. In fact, the more faithful and loyal a user is to Apple, the less likely it is they'll be able to enable advanced data protection. I just double-checked as I was preparing the notes on Sunday. I tried to enable it. I was provided with a list of six older but still in use by me Apple devices that would need to be running a newer edition of iOS or iPadOS than they're capable of

running. So ADP is a non-starter for me since I still use those older and still-working Apple devices every day.

Anyway, The Post continues, saying: "Technologists, some intelligence officers, and political supporters of encryption reacted strongly to the revelation after this story first appeared. Senator Ron Wyden, a Democrat on the Senate Intelligence Committee, said it was important for the United States to dissuade Britain. He said: 'Trump and American tech companies letting foreign governments secretly spy on Americans would be unconscionable and an unmitigated disaster for Americans' privacy and our national security.'

"Meredith Whittaker, of course who we know is the president of nonprofit encrypted messenger Signal, said: 'Using Technical Capability Notices to weaken encryption around the globe is a shocking move that will position the UK as a tech pariah, rather than a tech leader. If implemented, the directive will create a dangerous cybersecurity vulnerability in the nervous system of our global economy.'" Now, she didn't say they would pull out, but we know they would. She has previously said that when the EU was rattling their sabers similarly.

"Law enforcement authorities," writes The Post, "around the world have complained about increased use of encryption in communication modes beyond simple phone traffic, which in the United States can be monitored with a court's permission." And as we know, can also be monitored without the court's permission by China. "The UK and FBI in particular have said that encryption lets terrorists and child abusers hide more easily. Tech companies have pushed back, stressing a right to privacy in personal communication and arguing that backdoors for law enforcement are often exploited by criminals and can be abused by authoritarian regimes.

"Most electronic communication is encrypted to some degree as it passes through privately owned systems before reaching its destination. Usually such intermediaries as email providers and Internet access companies can obtain the plaintext if police ask. But an increasing number of tech offerings are encrypted end to end, meaning that no intermediary has access to the digital keys that would unlock the content. That includes Signal messages, Meta's WhatsApp, which as we know is based on Signal, and Messenger - WhatsApp and Messenger both from Meta - and of course Apple's iMessages and FaceTime calls. Often such content loses its end-to-end protection when it's backed up for storage in the cloud. That does not happen when Apple's Advanced Data Protection option is enabled.

"Apple has made privacy a selling point for its phones for years, a stance that was enhanced in 2016 when it successfully fought a U.S. order to unlock the iPhone of a dead terrorist in San Bernardino, California. It has since sought to compromise, such as by developing a plan to scan user devices for illegal material." I'll mention that again in a second. "That initiative was shelved after heated criticism by privacy advocates and security experts, who said it would turn the technology against customers in unpredictable ways.

"Google would be a bigger target for UK officials because it's made the backups for Android phones encrypted by default since 2018. Google spokesman Ed Fernandez declined to say whether any government had sought a backdoor, but implied none have been implemented. He said: 'Google cannot access Android end-to-end encrypted backup data, even with a legal order.' Meta also offers encrypted backups for WhatsApp. A spokesperson declined to comment on government requests but pointed to a transparency statement on its website saying that no backdoors or weakened architecture would be implemented. If the UK secures access to the encrypted data, other countries that have allowed encrypted storage, such as China, might be prompted

to demand equal backdoor access, potentially prompting Apple to withdraw the service rather than comply." And of course that's what everyone thinks they'll do.

"The battle over storage privacy escalated in Britain is not entirely unexpected. In 2022, UK officials condemned Apple's plans to introduce strong encryption for storage. A government spokesperson told the Guardian newspaper, referring specifically to child safety laws: 'End-to-end encryption cannot be allowed to hamper efforts to catch perpetrators of the most serious crimes.'

"After the Home Office gave Apple a draft of what would become a backdoor order, the company hinted to lawmakers and the public what might lie ahead. During a debate in Parliament over amendments to the Investigatory Powers Act, Apple warned last March that the law allowed the government to demand backdoors that could apply around the world. In a written submission, Apple stated: 'These provisions could be used to force a company like Apple, that would never build a backdoor into its products, to publicly withdraw critical security features from the UK market, depriving UK users of these protections.'

"Apple argued that when wielding the act against strong encryption would conflict with a ruling by the European Court of Human Rights, that any law requiring companies to produce end-to-end encrypted communications 'risks amounting to a requirement that providers of such services weaken the encryption mechanism for all users' and violates the European right to privacy."

Finally: "In the United States, decades of complaints from law enforcement about encryption have recently been sidelined by massive hacks by suspected Chinese government agents, who breached the biggest communications companies and listened in on calls at will. In a joint December press briefing on the case by FBI leaders, a Department of Homeland Security official urged Americans not to rely on standard phone service for privacy and to use encrypted services when possible." And we mentioned that at the time. "Also that month, the FBI, the NSA, and CISA joined in recommending dozens of steps to counter the Chinese hacking spree, including 'Ensure that traffic is end-to-end encrypted to the maximum extent possible.' Officials in Canada, New Zealand, and Australia endorsed the recommendations; those in the United Kingdom did not."

Okay. So The Washington Post's report correctly noted, and as we analyzed after its architecture was published, Apple has properly implemented true end-to-end encryption for every one of its cloud-based services where its use is feasible. As such, only the user's various iOS and iPadOS devices contain the key that's required to decrypt the contents of the data stored and shared in the cloud. Everything transiting to and from the cloud is, as we used to say, PIE - Pre-Internet Encrypted - and cannot possibly be accessed by anyone with access to either the data stored or in transit. The data can only be encrypted or decrypted on the user's device, and the key can never be removed from the user's device.

So we're back here once again, with the UK demanding something that none of the providers of secure messaging or secure storage will be willing to accommodate. But there's been a recent change that promises to provide the long sought-after solution to at least part of this problem, at least one of the reasons that everybody, like the bureaucrats and politicians are saying they need this, and that's for the children. And that's AI. Back in 1964, as part of a ruling about pornography, U.S. Supreme Court Justice Potter Stewart famously said: "I may not be able to define it, but I know it when I see it."

I see no reason why AI, functioning as an autonomous angel perched on every iOS user's shoulder, should not be able to stand in for Justice Stewart. This AI would not need to

contain the library of known CSAM - Child Sexual Abuse Material - the hashes for which users refused to have pre-loaded into their devices, feeling that this awful stuff was somehow in their phone. Instead, an AI would be trained to recognize such images. We know that Apple devices are already actively performing some of this "nanny" function. They are already empowered to warn their underage users when they may be about to send or receive and view any imagery that might be age-inappropriate for them. And this is all that any far more capable AI-enabled monitoring system would need to do.

What's significant is that it would not need to prevent the device from capturing and containing whatever content its user may wish to have. Parents can still take photos of their own kids in the bath. The system simply needs to filter out and prohibit the device's communication - its reception or transmission - of any such content that could potentially be subject to abuse. And once such filters are in place, there will be no need to gain access to anything stored in the cloud because there will be no way for anything abusive to leave or be received by any Apple device.

Given the history of government abuse of surveillance powers, many argue that the urgency to "protect the children" is just a smokescreen behind which lies a thirst for wider surveillance that could be turned, as it has been elsewhere, onto political rivals and other non-juveniles. So having companies like Apple, Signal, Meta, and others deploying local AI to lock down the content which their systems would refuse to send or receive short-circuits any governmental attempt at overreach.

And one of the best things about such solutions is that their effectiveness is so readily tested. Just present an AI-protected device with some test content that should not be communicated in order to verify that it's doing its job. So I really - I can see this, you know, the world is all abuzz about AI. We're understanding how capable it is. It seems easily possible that a local competent AI image recognition system could perform filtering functions on individual users' devices.

Leo: Well, yes, I guess. I mean, yeah. I think people - it wasn't merely that they didn't want the key, the hashes on there. I think they just don't like the idea of that kind of scanning going on in their phone.

Steve: Of any involvement of any kind.

Leo: Yeah, yeah. Maybe, you know, if it's a trade for that to encryption, I think it's obvious they want everything. This CSAM is just a pretext. They want everything.

Steve: Yeah, yeah. Also, as I mentioned at the top, France is doing something similar. An article appearing in Intelligence Online carried the headline "France Makes New Push for Backdoors Into Encrypted Messaging Apps." And the additional detail about that was behind a paywall. But also showing it said "French senators have passed an amendment paving the way for intelligence agencies to access backdoors into messaging apps such as WhatsApp, Signal, and Telegram." And presumably, you know, iMessage. What we believe is there are no such backdoors. So that would be requiring them to compel their creation. It's going to be interesting, Leo, to see what happens. You know, is Apple, what, going to say to the UK, well, we're going to not offer any encryption in the UK? But as we know, that's not what the UK is demanding. They demanding access to any user anywhere. Like, you know, demanding the end of encryption.

Leo: Right, basically, yeah. The thing that we don't know, and probably will never know, is because this was secret, I mean, the UK has not admitted to it, as you said, nobody - it's just it was a leaker. I would imagine they've also sent similar requests to Signal, WhatsApp, Google; right?

Steve: It just hasn't leaked; right. Because why would they target Apple? And Apple's not even the majority platform.

Leo: Right.

Steve: Google and Android are the larger platform.

Leo: Right. Why stop at Apple? So that, I mean, there will be no refuge except for doing something, a roll-your-own kind of a thing, if you really wanted them to end encryption.

Steve: And as we've said, if encryption is outlawed, only the outlaws will be using encryption.

Leo: Yeah, people with incentive. Because, frankly, very few people use Advanced Data Protection. You found one of the things that stopped me is you have to have everything up to date. But also you lose some capabilities, and there's this whole big risk of losing all of your data, too, if you forget your password. Most people don't use it. So, I mean, the people who are most motivated to use encryption, who are criminals, of course - well, no, not exclusively, but criminals are among those - are going to find ways. So this isn't going to have any effect. It's 1984 is what it is.

Steve: Yeah.

Leo: It's a bit depressing.

Steve: Okay. Another break, and then we're going to talk about Firefox 135.

Leo: Okay.

Steve: And a bunch of new features.

Leo: Yes, sir. I'll start downloading it right now. Steve?

Steve: So Firefox 135 was released one week ago, last Tuesday. And there's some interesting news about some new features. Despite having launched Firefox four days after last Tuesday, my Firefox was still on the previous 134 release. So I went to About Firefox, and that's how I saw that I was on 134, and it said, you know, update or upgrade or whatever, and I clicked a button, and it did that and restarted. I was first

greeted with a big page telling me that I'm now able to edit PDFs directly in Firefox, which may indeed come in handy.

But beyond that, Firefox Translations now supports more languages than ever. Pages in Simplified Chinese, Japanese, and Korean can now be translated; and Russian is now available as a target language for translating into. And in fact I used that a couple days ago for some Russian site that I went to when I was pursuing news for the podcast. And it came up unintelligible, but there was that little translation icon...

Leo: Oh, yeah, it's very useful, yeah.

Steve: ...at the right hand of the URL. I clicked it, and blink, it turned it into English. So actually I think some of the text that's in here is from the translation. So it's very handy.

Also, credit card autofill is now being rolled out gradually to all users globally. And as I mentioned, AI Chatbot access is now also being gradually rolled out. I already had it when I updated. To use it, you choose the AI Chatbot from the sidebar list of available sidebars, or you can go to Firefox Labs under the Settings page in order to find it. Then you choose which provider you want and so forth. I'll talk about that in more detail in a second.

Firefox also enforces certificate transparency, meaning that web servers must provide sufficient proof that their certificates were publicly disclosed before they will be trusted. And this only applies to servers using certificates that were issued by a certificate authority in Mozilla's Root CA Program. But that's all mainline certificate authorities. So that's just tightening up Firefox's public key certificate management.

Also good news, CRLite, which we've talked about, that's the Bloom filter-based CRL revocation system, is also now being gradually rolled out. So before long, from Firefox 135 on, we will have, as we discussed when we talked about this, Mozilla several times a day updating a master Bloom filter which our browsers will download, and then we will be doing browser-side revocation checking with very short delay, and no privacy concerns. Our browsers will not be reaching out to anybody asking whether the certificates that they're receiving from web servers are still valid.

Firefox now includes, they wrote, safeguards to prevent sites from abusing the history API by generating excessive history entries. I'm sure we've run across this. It bugs me when this happens. That makes navigating with the back and forward buttons difficult by deliberately cluttering up the history. You know, you go to a page, and it refers you to another page, but then the back arrow doesn't allow you to get back to where you came from. Sometimes you're able to hit back very quickly several times in order to get around that, but not always. So they've built that in so that only the history can no longer be inserted through JavaScript API without the user actually taking actions that create a breadcrumb history.

They also said that the "Do Not Track" checkbox has been removed from preferences. That's only because it's been incorporated into the global privacy control. So GPC is where that much stronger protection has been incorporated. And the "Copy Without Site Tracking" menu item was renamed "Copy Clean Link." Basically, if you're copying a link that has tracking crap in it, Firefox will remove that debris in order to give you a clean link. So it's now called Copy Clean Link rather than Copy Without Site Tracking.

And that's about half of - those were the most interesting half of the changes that are now in 135. Being a user myself, as we know, of ChatGPT, the idea of having it even more handy in my Firefox sidebar, where normally I always have Tree Style Tabs open

there, that's intriguing. As I said, I already have access to it. It'll be interesting to see what percentage of our listeners do. They're saying it's being rolled out, but it already came to me. It's CTRL+ALT+X is the shortcut which immediately jumps you to the AI chat in the sidebar. And at the moment Anthropic's Claude, ChatGPT, Google's Gemini, HuggingChat and Le Chat Mistral are the various AIs that are supported. You're able to choose among them and jump around them dynamically, as well.

So anyway, you can also, if you go to the Settings page under the hamburger menu icon in the upper-right, and then under Settings over on the left go to Firefox Labs, you're able to enable it and see if it's available on your browser, if you couldn't get to it in the sidebar. So anyway, bunch of cool things added to our favorite browser that, Leo, at least you and I use it, and I know that a lot of our listeners do, too.

The United States National Security Agency, our NSA, in coordination with our four partner countries which together form the Five Eyes alliance, has just released the latest guidance on securing network edge devices. I'm just going to share their joint announcement, which is relatively short, but I know that many of our listeners have frontline responsibility in their enterprises with a great many necessarily exposed devices on the edge. Meaning, you know, the network edge, typically the edge where the Internet connects to the enterprise.

So the NSA.gov site's release of this, in coordination, it was dated from Fort Meade, Maryland, and they said: "The National Security Agency has joined the Australian Signals Directorate's Australian Cyber Security Centre, the Canadian Centre for Cyber Security, and others to release three Cybersecurity Information Sheets." Of course everything is an acronym with these guys, so they're CSIs, Cybersecurity - it ought to be CISEs, Cybersecurity Information Sheets - that highlight critically important mitigation strategies for securing edge devices, including firewalls, routers, and virtual private network gateways.

Collectively, these reports are "Mitigation Strategies for Edge Devices," and the first sheet is the Executive Guidance. The second is "Mitigation Strategies for Edge Devices: Practitioners Guidance." And then the other is "Security Considerations for Edge Devices." They said they provide high-level summary. So I've got links in the show notes to the announcement from the NSA of this, and in the announcement are the links to each of those three reports.

And, you know, the executive guidance is a broad overview. You know, know the edge, procure secure-by-design devices, apply hardening guidance and so forth, you know, sort of basic stuff. The security guidance is much more detailed and very useful. But it occurred to me that, for our listeners, it's always useful to have a checklist; right? Just to go through and say, yup, took care of that. Yup, considered that. Yup, considered that. And that's nice for covering one's butt. If anything does happen, you're able to say, well, you know, we're in full compliance with the NSA's latest guidance. And there's also a sheet you can give to your boss and say, look, boss, we need to buy some stuff here because, you know, our stuff won't do what the NSA is telling us we need to do. So, useful info.

I mentioned Netgear at the top. Anyone having a recent Netgear WiFi-6 access point or Netgear Nighthawk gaming router should be very sure that you're running the latest recently released, as of last week, firmware. Make sure now. Three Netgear WiFi-6 devices, the models WAX214v2, also that same WAX206 and WAX220, those three models, until and unless updated, all contain highly critical CVSS 9.6 authentication bypass vulnerabilities. And we know what that means. If there's anything exposed to the Internet, there's now a way for bad guys to get in. And as I said, I already know that as a follower of this podcast you would never enable any Internet-facing remote management capabilities.

Leo: Never. Never. No.

Steve: No. But it's also human to assume that it could never happen to you. So please make sure that you're running the latest firmware as of last week. And better yet, arrange to never be vulnerable in the first place by not opening any of those sorts of ports. So those three WiFi routers were vulnerable to a now-patched authentication bypass with that 9.6 out of 10.

But three other Netgear Nighthawk gaming routers rated an even higher CVSS score of 9.8 for their unauthenticated remote code execution vulnerabilities. The three affected routers are the XR500, the XR1000, and the XR1000v2. They are all Nighthawk WiFi 6 Pro Gaming Routers. If any of those numbers sound familiar, and especially if you or someone you know may have been unable to resist the temptation of enabling any sort of remote access, you'll want to update them to the latest firmware immediately. I saw no reports of this being a zero-day. As far as I know, the vulnerability was responsibly reported to Netgear. But we also know that, once it is known that these problems exist, bad guys can reverse engineer the firmware in the unpatched routers, figure out how to get in, and then start attacking. So there is a window here. You want to make sure that you're not vulnerable within that time period.

Okay. Sysinternals. There was a surprising bit of news involving the much beloved Sysinternals tools. As many of our listeners know, they were a collection, and still are, of truly unique and powerful utilities that were originally created by Mark Russinovich and Bryce Cogswell. Their little Texas-based company was purchased lock, stock, and barrel by Microsoft back in 2006, much to many people's chagrin, since everyone was quite worried at the time that it might spell the end of that fabulous and really irreplaceable tool set. Fortunately, that didn't happen, and the tools remain available today from Microsoft and are still being maintained and upgraded.

Which makes this news of a recent discovery all the more curious and troubling. A software engineer by the name of Raik Schneider has reported that he has discovered DLL hijacking bugs in the Sysinternals tools. Oh, in fact, it's this guy's page. It's written in German, and it was Firefox's built-in translator that allowed me to turn it into English. So the curious and troubling part is that Microsoft has done nothing about these problems. They remain unpatched. And, worse, their existence is now public and widely known, even after a 90-day responsible disclosure window.

So Raik's detailed public disclosure reads, and this is just the beginning of it, he said: "I have identified and verified critical vulnerabilities in almost all Sysinternals tools and presented the background and attack in a video. A summary of the weak spot and the link to the video can be found here in this blog post. These tools, developed by Microsoft - and actually originally Sysinternals of course - are widely used in IT administration and are often used for analysis and troubleshooting. The vulnerability demonstrated in the video affects numerous applications of the suite and allows attackers to use DLL injection to inject and execute defective code."

And, now, okay, that may be part of the translation. We know it could be, not defective, but malicious code. And he said: "Now that more than 90 days have passed since the initial disclosure to Microsoft, it's time to talk about it." And then he goes on to do so. I have a link to his posting in German in the show notes. And if you've got a translator built into your browser and don't speak German, then it'll do a good job of translating it into English for you.

Leo: Actually, my translation, and I'm not sure where it came from, says "malicious code."

Steve: Ah, interesting. Okay.

Leo: Yeah. So this is Arc. So I don't know what translator it's using.

Steve: Oh, interesting.

Leo: Yeah, probably not Google.

Steve: Okay. So the problem is a well-known and common problem with Windows DLLs where, among many problems, DLLs made sense back when we had 128MB Windows 2 computers because it was a way of sharing code. And the idea would be that, rather than various applications all needing to bring their own code along, not only because we had applications that were sharing 20MB hard drives or floppies, but because there wasn't much RAM. So you didn't want - so you just wanted to be able to share these libraries. Great idea back then. Today it's pure legacy. It absolutely makes no sense whatsoever. But there's never been a point in time where Microsoft could break this. So we still have it today.

So what happens is the Windows executable file loader, when it's loading an executable file, is able to, in the executable, the executable declares the DLLs that it's reliant upon, the system code DLLs that it needs. And so the executable file loader loads those for the executable so that they're there and linked up to it and ready to go. It first looks in the application's own directory, that is, where the EXE is being run from. And this behavior was originally deliberate since it allowed applications to bring along their own more recent or maybe even older versions of DLLs. This is where the whole DLL thing began to fall apart because they would then be loaded and used preferentially over whatever same-named DLLs the system might already or might not have.

The problem is, that convenience feature can be readily abused. In the case of the Sysinternals executables, they're not relying upon any of their own DLLs. This is actually one of the things that makes them so nice is that they're single executables that just get their jobs done. Very clean. But like all Windows applications, they DO rely heavily upon many system DLLs. But rather than insisting that the system DLLs they require be loaded from within the system's own protected directories, as they should, the Sysinternals apps use the default behavior, where Windows will first look inside the app's own directory. And this enables the exploit. Bad guys can place a DLL that's named the same as a system DLL in Sysinternal's execution directory, and it will be loaded instead of the intended system DLL.

This flaw has been widely picked up and reported by the tech press over the past few days. The reporting notes that many of the Sysinternals utilities prioritize DLL loading from untrusted paths such as the current working directory or network paths, before looking in secure system directories for their DLLs. One piece of this reporting wrote: "The vulnerability was responsibly disclosed to Microsoft on October 28th, 2024. However, Microsoft classified it as a 'defense-in-depth' issue rather than a critical flaw. This classification implies that mitigation relies on secure usage practices rather than addressing it as a fundamental security defect. While Microsoft emphasizes running executables from local program directories, researchers argue that network drives where

the current working directory becomes the application's execution path pose significant risks, as indeed they do."

So what's most significant here to me is the breadth of press coverage and reporting that this news has generated. I mean, this got picked up because Sysinternals is so popular, this got picked up everywhere.

Leo: Yeah, everybody uses it, yeah.

Steve: Yes. So we've seen Microsoft respond when sufficient noise is made. We saw how quickly they backpedaled on the first release of their Copilot+ "Recall" screen scraper. So I would imagine that the amount of bad press that is being generated here will result in someone's attention being pointed at updating all of the vulnerable Sysinternal tools. I suspect Microsoft is regretting that they blew this off and said, oh, it's not our problem. You just have to be careful how you use them. It's like, okay, good luck with that. Unfortunately, there's no update mechanism for the bazillion copies of Sysinternals tools that have already been downloaded and are deployed. They will never be updated.

Leo: Oh, interesting.

Steve: Unless they're manually replaced. They all have this behavior.

Leo: Yikes.

Steve: Yeah.

Leo: Okay.

Steve: This creates an enduring opportunity for exploitation.

Leo: What is a defense-in-depth issue? What does that mean? That just you should be careful.

Steve: Yeah. Exactly. This is like...

Leo: You should have done a better job.

Steve: Yeah, exactly. Like, well, yes. But, you know, these are advanced sleuthing tools. So you shouldn't leave them around on computers where they could be exploited.

Leo: Yeah.

Steve: Thanks, but everybody does.

Leo: Yeah. I imagine it's a custom DLL that the Sysinternals run. In fact, they probably have a common DLL.

Steve: No. No.

Leo: No?

Steve: It's like kernel32.dll needs to get loaded.

Leo: Oh, it should definitely be getting that from the secure...

Steve: Exactly. And they don't. They don't. So someone malicious names their malicious code kernel32.dll, puts it where the Sysinternals tool is, and that's the one that gets loaded.

Leo: Isn't this a widespread problem, though, in Microsoft?

Steve: Yeah. Yeah. I mean, it requires overriding Windows standard default behavior, which they can't change because it will break things that depend upon it.

Leo: Right. They didn't used to have a secure place to store those DLLs, actually.

Steve: Security was never a consideration. If you've got a Windows 2 machine with floppy disks, what security?

Leo: Security with a - what are you pretending?

Steve: Why not let a Windows metafile execute code in the image because that might come in handy. And that's what they did in the beginning.

Leo: You raise an excellent point, though. I think that DLLs are just running on inertia. There's no - you don't need it anymore.

Steve: There was never a time when they could afford to break this. I mean, you know...

Leo: I mean, on Linux you have static linked executables. I mean, that's - they're bigger because of it. But then you don't have that problem. You don't have libraries, and you don't have the DLL hell that you get with Windows.

Steve: Right.

Leo: With conflicting DLL versions and so forth. Hmm.

Steve: Yeah.

Leo: Maybe it's time to think about getting rid of those. Just don't do it anymore.

Steve: Maybe go VM happy and execute each EXE in its own VM.

Leo: That was the plan.

Steve: Yeah, I know. It's a mess. Okay. Google removes the ban on using AI for harm. What? Last Tuesday, Wired covered an interesting change in Google's policies regarding the conduct and use of its AI. Wired's headline was "Google Lifts a Ban on Using Its AI for Weapons and Surveillance."

Leo: Well, it's about time.

Steve: That's right.

Leo: This is the - when you talk about the existential threat of AI, this is the first thing that leaps into my mind. Right? Don't have autonomous nuclear weapons.

Steve: Yes. The tag line in Wired's coverage said: "Google published principles in 2018 barring its AI technology" - such as it was - "from being used for sensitive purposes. Weeks into President Donald Trump's second term, those guidelines are being overhauled." Okay, I have no idea why Wired referred to our current president's administration since there's no reason I can see to believe that there's any connection between the two.

Here's what Wired wrote. They said: "Google announced Tuesday that it is overhauling the principles governing how it uses artificial intelligence and other advanced technology. The company removed language promising not to pursue 'technologies that cause or are likely to cause overall harm,' and 'weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people.' And finally, 'technologies that gather or use information for surveillance violating internationally accepted norms.' And finally, 'Technologies whose purpose contravenes widely accepted principles of international law and human rights.' All that language was taken out. The changes were disclosed in a note appended to the top of a 2018 blog post unveiling the guidelines, saying, 'We've made updates to our AI principles. Visit AI.google for the latest,' the note reads.

So in a blog post on Tuesday a pair of Google executives cited the increasingly widespread use of AI, evolving standards and geopolitical battles over AI as the backdrop to why Google's principles needed to be overhauled.

Wired wrote: "Google first published the principles in 2018 as it moved to quell internal protests over the company's decision to work on a U.S. military drone program. In response, it declined to renew the government contract and also announced a set of

principles to guide future uses of its advanced technologies such as artificial intelligence. Among other measures, the principles stated Google would not develop weapons, certain surveillance systems, or technologies that undermine human rights. But in an announcement on Tuesday, Google did away with those commitments. The new web page no longer lists a set of banned uses for Google's AI initiatives. Instead, the revised document offers Google more room to pursue potentially sensitive use cases."

It states: "Google will implement appropriate human oversight, due diligence, and feedback mechanisms to align with user goals, social responsibility, and widely accepted principles of international law and human rights."

Google also now says it will work to "mitigate unintended or harmful outcomes." Which, okay, still says some of the same things, though maybe a little less pointedly. James Manyika, Google senior vice president for research, technology, and society, was quoted - so this guy a Google person: "We believe democracies should lead in AI development, guided by core values like freedom, equality, and respect for human rights." Right. And Demis Hassabis, the CEO of Google's esteemed AI research lab DeepMind, said: "And we believe that companies, governments, and organizations sharing these values should work together to create AI that protects people, promotes global growth, and supports national security."

They added that Google will continue to focus on AI programs "that align with our mission, our scientific focus, and our areas of expertise, and stay consistent with widely accepted principles of international law and human rights." At the same time, multiple Google employees expressed concern about the changes in conversations with Wired.

Okay. Well, my own feeling is that we should not read much into this, and I guess I salute Google for being upfront about it. I mean, they're not hiding at all behind the fact that they've changed their wording on this. These guidelines were first created seven years ago, back in 2018. Seven years in AI timeframe, you know, is Jurassic. The world of AI has obviously been dramatically transformed since then, and I suspect that this is just Google being upfront about needing to operate on a level playing field alongside everyone else. They could have left that language there and ignored it, if necessary. They're not saying that they're going to proactively "do bad." They're just saying that they're going to abide by the same rules as everyone else. So okay.

Leo: All right. Let's talk about - I've been, by the way, quoting last week's episode all week long with the bottom line being there is really no AI that hasn't been jailbroken.

Steve: No.

Leo: AI safety is an illusion, basically.

Steve: And my intuition is that we're going to have a hard time putting guardrails around AI. It, I mean, we were surprised when this worked at all. There's still big questions, like, surrounding how does it work at all.

Leo: Yes. We don't really even know how it works.

Steve: No. So, you know, it's like, okay, if you don't know how it works, how are you going to tell it not to talk about some things that it knows. I mean...

Leo: The old hacker creed was information wants to be free. AI wants to be free. It wants to help you. It wants to tell you what you want to know. Pretty hard to stop it.

Steve: So, yes. Following up on last week's look at the relative weakness of the DeepSeek AI model resistance to jailbreaking, we have a post on LinkedIn by Eran Shimony whose title is "Principal Vulnerability Researcher at CyberArk." Eran's post reads: "OpenAI recently released the o3 family of models" - right, that's the top end best there is right now - "showcasing significant advancements in reasoning and runtime inference. Given its expected widespread use in development, ensuring it does not generate malicious code is crucial. OpenAI has strengthened its security guardrails, mitigating many previous jailbreak techniques." You know, and Leo, to your point, you called it Whac-A-Mole last week, and that's exactly right.

Leo: Yeah, yeah.

Steve: I mean, it's like, oh, how about try this. Oops. Okay, let's go fix that. How about this? Oops. Oh, good, let's go fix that. Doesn't feel solid.

Leo: No.

Steve: He said: "However, using our open source tool, FuzzyAI, we successfully jailbroke o3" - get this - "extracting detailed instructions on injecting code into lsass.exe, including a breakdown of the obstacles involved, ultimately leading to functional exploit code."

Leo: Oh, my god.

Steve: Okay, now, lsass.exe always shows up in any list of running Windows processes.

Leo: I know. I've googled it, saying what the hell is this?

Steve: Yes. LSASS stands for Local Security Authority Subsystem Service. It is the Security God of Windows because it's the Windows process that manages user authentication and all security policies. Being able to inject attack code into that process would create the mother of all privilege escalation and restriction bypasses. And these guys tricked ChatGPT's latest and most powerful code-generating o3 model to write the code to do just that.

Leo: Which on the one hand is extremely impressive. Right? Like, wow, it wrote the code, wow. On the other hand, it's depressing.

Steve: It is impressive. But it's also very worrisome.

Leo: Yes.

Steve: Right. He said: "While AI security is improving, our findings indicate that vulnerabilities still exist, highlighting the need for further safeguards. We've opened a Discord community for our open-source tool. You're welcome to join."

Leo: I just tried to join it. It's gone.

Steve: Really.

Leo: Yeah, I'm not sure if they got booted or...

Steve: I got an invite when I went there yesterday, so...

Leo: Is it there? Were you able to get in?

Steve: I did yesterday.

Leo: Oh, okay. Well, maybe...

Steve: Using that link.

Leo: Oh, maybe there's something I'm doing wrong, then. I'll try again.

Steve: It's waiting to come up. And, yup, I got - oh, no, invalid invite.

Leo: Yeah, that's what I got.

Steve: Invite may be expired, or you might not have permission to join. Yup.

Leo: Right. That's what I got.

Steve: So probably - and it was a LinkedIn obscured link, which I de-LinkedIn-a-fied. But you can also go to his GitHub page. It's [GitHub.com/cyberark](https://github.com/cyberark), and then look for the project FuzzyAI. And in fact it might be that if you go there you'll find the Discord invite.

Anyway, the FuzzyAI GitHub page says: "The FuzzyAI Fuzzer is a powerful tool for automated LLM fuzzing. It's designed to help developers and security researchers identify jailbreaks and mitigate potential security vulnerabilities in their LLM APIs. It features Comprehensive Fuzzing Techniques: Leverage mutation-based, generation-based, and intelligent fuzzing. Built-in Input Generation: Generate valid and invalid inputs for exhaustive testing. Seamless Integration: Easily incorporate into your development and

testing workflows. And Extensible Architecture: Customize and expand the fuzzer to meet your unique requirements. And it supports Anthropic's Claude; OpenAI's GPT-4o; Gemini's Gemini Pro; Azure's GPT-4 and GPT-3.5 Turbo; Bedrock that has Claude; AI21's Jamba; DeepSeek's both V3 and V1; and Ollama, both LLaMA and Dolphin-LLaMA."

So this sort of research and experimentation is exactly what is needed, so a big Bravo to CyberArk. And Leo, it feels to me as though the problem is inherently intractable. And I'm sure this is a major source of anxiety for AI developers. The problem is that you're not going to get a clean edge. You're not going to be able to create a clean boundary.

So that if - and my point is, in order to prevent these things from answering questions you don't want them to, from generating code you don't want them to, you're going to have to so restrict them that they are no longer able to answer questions you do want them to, and generate code you would like them to be able to do because there just isn't, you know, like what's the boundary between something malicious and not? It's sort of your view; right? It's, you know, one person's malicious code is another person's requirement for solving a problem in IT.

Leo: Right. Right.

Steve: So, you know, we're dealing with fuzzy definitions.

Leo: It's very subjective, absolutely.

Steve: And when your definitions are fuzzy, how can you expect the AI to make that determination?

Leo: Right. Right. All you can do is give it a list of words and things. You know? I mean, that's really all they're doing is saying, if somebody says Tiananmen Square, make sure you don't say anything about that. And that's an infinite list. You can't ever get everything.

Steve: Right. And in fact that...

Leo: And you can get around it. What is fuzzing? So fuzzing in this context, is it the same as fuzzing in other exploit generating?

Steve: Yeah, it's basically just trying to confuse it. Well, okay, so...

Leo: It's almost like randomized prompts.

Steve: It's not the same thing, inasmuch as, you know, fuzzing data to a port is very specific. But yes, it is, you know, feeding gibberish in and seeing what comes out. And like we know that the models grow the context over time, and that that context is one of the things that gives them the power that they have for - it's what makes the model interactive and allows you to say, oh, I'm sorry, I didn't explain what I wanted correctly. I meant more like this. And allows you to set up scenarios that allow models to be

tricked. And so the more we automate this and the more crap we throw at the wall, the more we're able to see whether we're able to get an answer that the designers didn't intend the model to produce. So it'll be useful for them, for the designers, as well. It's a mess, Leo.

Leo: It also can do it at speed.

Steve: Yes.

Leo: Which is a big advantage.

Steve: Yes.

Leo: You can throw a lot of stuff at it.

Steve: As long as you're able to afford the API cost.

Leo: Right.

Steve: You know, it's not going to be cheap to run lots of deep inferences as this fuzzing would require. But costs are going to come down, too.

Leo: Right.

Steve: Next page is something that I am very excited to share.

Leo: I like this. Oh, I like this. Okay.

Steve: I saw it for the first time myself yesterday evening. It is a screenshot of GRC's DNS Benchmark, which is the first ever simultaneous multiprotocol Benchmark of Name Servers showing DNS over HTTPS, DNS Over TLS, IPv4 and IPv6 Name Servers, all being benchmarked at once and with their performance compared against each other.

Leo: Wow.

Steve: And the preliminary results are interesting. That fastest of all is, at the very top there, is NextDNS's DNS over HTTPS name server.

Leo: No. That's what I use. Oh.

Steve: But you have to be using HTTPS, so you have to be using DNS over HTTPS.

Leo: Oh, okay, okay.

Steve: So you would be - you'd need to configure your web browser to do that. My guess is it's fastest because it's not being heavily used yet.

Leo: Right. No one's using it yet.

Steve: Uh-huh.

Leo: Yeah.

Steve: And then in the number two place, you'll notice under the bars at the top it says "determining ownership." That's still old code. The Benchmark always used to just resolve IP addresses, so there's the system called Sender Base that allows you to give it an IP and look up the owner of that IP space. Well, that's not widely supported for URL-based name servers, which is what DOH and DOT are.

Leo: Ah.

Steve: Anyway, so my point is that what will be shown shortly, I mean, this just came to life last night.

Leo: This is [crosstalk].

Steve: Yes.

Leo: And we should mention that these results are local to you. That's why everybody needs to run their own.

Steve: Yes, that's exactly, exactly right. It is from my location in Southern California. That's what I saw. And in fact what isn't there is normally these bars would have been squished way down by the other name servers that were so slow by comparison. I deleted them in order so that only - so that you could see at the bottom is that one green bar is the one that is the slowest of all.

Leo: Oh, yeah.

Steve: That it's what set the scale for everything else. Anyway, the second one from the top is Quad 9's.

Leo: I see Quad 1 and Quad 9, yeah.

Steve: Yeah. And Quad 9's is that second fastest, DNS over TLS. So, yeah, so from my position in Southern California, that's what I saw. I've got more work to do on the UI. But I will be producing the fifth release of this for testing by our gang probably in the next few days.

Leo: Is this multithreaded? It must be.

Steve: Oh, my god. It's crazy multithreaded. I mean, everything is running at once.

Leo: Wow, it's so cool.

Steve: It is. It is really.

Leo: How many processors do you have in this machine?

Steve: It'll run multithreaded. Remember, everything's in assembler. It's still only a couple of hundred K because, you know, it is super efficient. And actually doing DNS queries is not time-consuming.

Leo: Sure.

Steve: It's just sending a short packet out and then timing how long it takes for it to come back.

Leo: Yeah.

Steve: So, yeah. It is a massively parallel application.

Leo: Oh, that's awesome.

Steve: But it's starting to come to life. So anyway, I'm very, very happy to be able to share that and show it. And I'll get it done.

Leo: Yay.

Steve: Okay. We're going to now talk about the hidden fact of ransomware attacks in K-12 schools in the U.S. I don't know whether or not it would come as a surprise that hiding school cyberattacks is a thing. You know? It might come as a surprise that it's actually a job description.

Leo: Oh, really.

Steve: It is. There are people whose job description is hiding school cyberattacks, and they're being paid to do it. So exactly one week ago, last Tuesday, the website of an organization called "The 74" published an eye-opening piece of investigative journalism that I knew would make a terrific topic for the podcast. As I said at the top of the show, "74" stands for 74 million, which is the number of American school-age children being educated from kindergarten through high school in the U.S.

The 74's code of reporting ethics states: "The 74 is a nonprofit, nonpartisan national news organization covering K-12 education. The organization's mission is to spotlight innovative thinking and models that are helping students succeed, to cover and analyze education policy and politics, and to use journalism to challenge the conditions that deny too many children access to a quality education. The 74 is committed to reporting stories without fear or favor about what is working well for students and families, and to expose and hold accountable the systems that are failing them." And I took some time browsing around there, and it looks like a neat organization.

So last Tuesday this group published a story titled "Kept in the Dark Meet the Hired Guns Who Make Sure School Cyberattacks Stay Hidden." Here's what they reported. They said: "An investigation by The 74 shows that while schools have faced an onslaught of cyberattacks since the pandemic disrupted education nationwide five years ago, district leaders across the country have employed a pervasive pattern of obfuscation that leaves the real victims in the dark.

"An in-depth analysis chronicling more than 300 school cyberattacks over the past five years reveals the degree to which school leaders in virtually every state repeatedly provide false assurances to students, parents, and staff about the security of their sensitive information. At the same time, consultants and lawyers steer 'privileged investigations,' which keep key details hidden from the public. In more than two dozen cases, educators were forced to backtrack months and in some cases more than a year later after telling their communities that sensitive information, which included, in part, special education accommodations, mental health challenges, and student sexual misconduct reports that had not been exposed. While many school officials offered evasive storylines, others refused to acknowledge basic details about cyberattacks and their effects on individuals, even after the hackers made student and teacher information public.

"The hollowness in schools' messaging is no coincidence because the first people alerted following a school cyberattack are generally neither the public nor the police. District incident response plans place insurance companies and their phalanxes of privacy attorneys first. They take over the response, with a focus on limiting schools' exposure to lawsuits by aggrieved parents or employees. Attorneys, often employed by just a handful of law firms, dubbed breach mills by one law professor for their massive caseloads, hire forensic cyber analysts, crisis communicators, and ransom negotiators on schools' behalf, immediately placing the discussions under the shield of attorney-client privilege. Data privacy compliance is a growth industry for these specialized lawyers, who work to control the narrative.

"As a result, students, families, and district employees whose personal data was published online from their financial and medical information to traumatic events in young people's lives are left clueless about their exposure and risks to identity theft, fraud, and other forms of online exploitation. Told sooner, they could have taken steps to protect themselves. Similarly, the public is often unaware when school officials quietly agree in closed-door meetings to pay the cyber gangs' ransom demands in order to recover their files and unlock their computer systems. Research suggests that the surge in incidents has been fueled, at least in part, by insurers' willingness to pay. Hackers themselves have stated that when a target carries cyber insurance, ransom payments are all but guaranteed.

"In 2023, there were 121 ransomware attacks on U.S. K-12 schools and colleges, according to Comparitech, a consumer-focused cybersecurity website whose researchers acknowledge that the number is an undercount. For the same year, an analysis by Malwarebytes reported 265 ransomware attacks against the education sector globally in 2023, a 70% year-over-year surge, making it 'the worst ransomware year on record for education.' Daniel Schwarcz, a University of Minnesota law professor, wrote a 2023 report for the Harvard Journal of Law & Technology criticizing the confidentiality and doublespeak that shroud school cyberattacks as soon as the lawyers often called breach coaches arrive on the scene. Schwarcz told The 74: 'There's a fine line between misleading and, you know, technically accurate. What breach coaches try to do is push right up to that line, and sometimes they cross it.'

"The 74's investigation into the behind-the-scenes decision-making that undermines what, when, and how school districts reveal cyberattacks is based on thousands of documents obtained through public records requests from more than two dozen districts and school spending data that links to the law firms, ransomware negotiators, and other consultants hired to run district responses." All of this otherwise kept off the books and private, of course. It also includes an analysis of millions of stolen school district records uploaded to cyber gangs' leak sites. Some of students' most sensitive information lives indefinitely on the dark web, while other personal data can be found online with little more than a Google search, even as school districts deny that their records were stolen and cyber thieves boast about their latest score.

"The 74 tracked news accounts and relied on its own investigative reporting in Los Angeles; Minneapolis; Providence, Rhode Island; and Louisiana's St. Landry Parish, which uncovered the full extent of school data breaches, countering school officials' false or misleading assertions. As a result, district administrators had to publicly acknowledge data breaches to victims or state regulators for the first time, or retract denials about the leak of thousands of students' detailed psychological records.

"In many instances, The 74 relied on mandated data breach notices that certain states, like Maine and California, report publicly. The notices were sent to residents in these states when their personal information was compromised, including numerous times when the school that suffered the cyberattack was hundreds, and in some cases thousands, of miles away. The legally required notices repeatedly revealed discrepancies between what school districts told the public early on and what they later disclosed to regulators after extensive delays. Some schools, meanwhile, failed to disclose data breaches, which they are required to do under state privacy laws. And for dozens of other schools, The 74 could find no information at all about alleged school cyberattacks uncovered by its reporting, suggesting they had never before been reported or publicly acknowledged by local school officials.

"Education leaders who responded to The 74's investigation results said any lack of transparency on their part was centered on preserving the integrity of the investigation [uh-huh], not self-protection. School officials in Reeds Spring, Missouri, said: 'When we respond to potential security incidents, our focus is on accuracy and compliance, not downplaying the severity.' Those at Florida's River City Science Academy said the school 'acted promptly to assess and mitigate risks, always prioritizing the safety and privacy of our students, families, and employees.' In Hillsborough County Public Schools in Tampa, Florida, administrators in the nation's seventh-largest district said they notified student breach victims 'by email, mail, and a telephone call' and 'set up a special hotline for affected families to answer questions.'"

Hackers have exploited officials' public statements on cyberattacks to strengthen their bargaining position, a reality educators cite when endorsing secrecy during ransom negotiations. Doug Levin, who advises school districts after cyberattacks and is the co-founder and national director of the nonprofit K12 Security Information eXchange said:

"But those negotiations do not go on forever. A lot of these districts come out saying, 'We're not paying,'" the ransom. In which case the negotiation is over, and they then need to come clean. The paid professionals who arrive in the wake of a school cyberattack are held up to the public as an encouraging sign. School leaders announce reassuringly that specialists were promptly hired to assess the damage, mitigate the harm, and restore their systems to working order.

This promise of control and normality is particularly potent when cyberattacks suddenly cripple school systems, forcing them to shut down for days and disable online learning tools. News reports are fond of saying that educators were forced to teach students "the old-fashioned way, with books and paper." But what isn't as apparent to students, parents, and district employees is that these individuals are not there to protect them, but to protect schools from them.

And Leo, let's take our final break, and then I'm going to finish with this and then discuss it a little bit.

Leo: Okay, good. It's a little upsetting.

Steve: Yeah. It is. Going on behind the scenes and, you know, deliberately obscured.

Leo: Yeah.

Steve: So when the Medusa ransomware gang attacked Minneapolis Public Schools in February of '23, it stole reams of sensitive information and demanded \$4.5 million in bitcoin in exchange for not leaking it. District officials had a lawyer at Mullen Coughlin notify the FBI. So at the same time officials were not acknowledging publicly that they had been hit by a ransomware attack, their attorneys were telling federal law enforcement that the district immediately determined its network had been encrypted, promptly identified Medusa as the culprit, and within a day had its "third-party forensic investigation firm" communicating with the gang regarding the ransom.

Mullen Coughlin then told the FBI that it was leading "a privileged investigation" into the attack and, at the school district's request, "all questions, communication, and requests in connection with this notification should be directed" to the law firm. Mullen Coughlin did not respond to requests for comment. Minneapolis school officials would wait seven months before notifying more than 100,000 people that their sensitive files were exposed, including documents detailing campus rape cases, child abuse inquiries, student mental health crises, and suspension reports. As of December 1st, all schools in Minnesota are now required to report cyberattacks to the state, but that information will be anonymous and not shared with the public.

One district took such a hands-off approach, leaving cyberattack recovery to the consultants' discretion, that they were left out of the loop and forced to later issue an apology. When an April 2023 letter to Camden educators arrived 13 months after a ransomware attack, it caused alarm. An administrator had to assure employees that the New Jersey district wasn't the target of a second attack. The letter was about the one more than a year ago. The attorneys had sent out notices after a significant delay and without the school's knowledge.

Other school leaders said when they were in the throes of a full-blown cyber crisis and ill-equipped to fight off cybercriminals on their own, law enforcement was not of much use, and insurers and outside consultants were often their best option. Don Ringelestein, the

executive director of technology at the Yorkville, Illinois school district said: "In terms of how law enforcement can help you out, there's really not a whole lot that can be done, to be honest." When the district was hit by a cyberattack prior to the pandemic, he said, a report to the FBI went nowhere. Instead, district administrators turned to their insurance company, which connected them to a breach coach, who then led all aspects of the incident response under attorney-client privilege.

Northern Bedford County Schools Superintendent Todd Beatty said the Pennsylvania district contacted the CISA to report a July 2024 attack, but "The problem is there's not enough funding and personnel for them to be able to be responsive to incidents." And too many incidents. Meanwhile, John VanWagoner, the Schools Superintendent in Traverse City, Michigan, claims insurance companies and third-party lawyers often leave district officials in the dark, too. Their insurance company presented school officials with the choice of several cybersecurity firms they could hire to recover from a March 2024 attack, VanWagoner said, but he didn't know where to go to vet if they were any good or not. He said it had been a community member, not a paid consultant, who first alerted district officials to the extent of the massive breach that forced school closures and involved 1.2TB of stolen data.

Breach notices and other incident response records obtained by The 74 show that a small group of law firms play an outsized role in school cyberattack recovery efforts throughout the country. Among them is McDonald Hopkins, where Michigan attorney Dominic Paluzzi co-chairs a 52-lawyer data privacy and cybersecurity practice. Some call him a "breach coach." He calls himself a "quarterback." After establishing attorney-client privilege, Paluzzi and his team call in outside agencies covered by a district's cyber insurance policy including forensic analysts, negotiators, public relations firms, data miners, notification vendors, credit-monitoring providers, and call centers. Yeah. And who pays for this? The taxpayer. Across all industries, the cybersecurity practice handled 2,300 incidents in 2023, 17% of which involved the education sector - which, Paluzzi noted, is not quite "always the best when it comes to the latest protections."

When asked why districts' initial response is often to deny the existence of a data breach, Paluzzi said, "Well, it takes time to understand whether an event rises to the level that would legally require disclosure and notification." Paluzzi said: "It's not the time to make assumptions, to say, 'We think this data has been compromised,' until we know that. If we start making assumptions, that starts our clock on legally mandated disclosure notices. We're going to have been in violation of a lot of the laws, and so what we say and when we say it are equally important." Which is why there are so many jokes about attorneys, of course. In other words, finessing the system.

They said, you know: "Once we've acknowledged that a breach has occurred, notification requirement clocks start ticking. So the longer we wait to acknowledge, apparently even to themselves, that anything more serious than an 'incident' is being investigated, the better." He said: "In the early stage, lawyers are trying to protect their client and avoid making any statements they would later have to later retract or correct." Uh-huh.

Paluzzi said: "While it often looks a bit canned and formulaic, it's often because we just don't know, and we're doing so many things. We're trying to get it contained, ensure the threat actor is not in our environment, and get up and running so we can continue with school and classes. And then we shift to whatever data is potentially out there and compromised." A data breach is confirmed, he said, only after "a full forensic review," a process that can take up to a year, and often only after it's completed are breaches disclosed and victims notified.

He said: "We run through not only the forensics, but through the data mining and document review effort. By doing that last part, we are able to actually pinpoint for John Smith that it was his Social Security number, right; and Jane Doe, that it's your medical

information," he said. "We try in most cases to get to that level of specificity, and our letters are very specific." So it sounds like a lot of billable hours, to me. Makes you sort of wonder whether the cure is more, you know, is worse than the disease.

"According to," they wrote, "a 2023 blog post by attorneys at the firm Troutman Pepper Locke, targets that respond to cyberattacks without the help of a breach coach often fail to notify victims and, in some cases, provide more information than they should. When entities over-notify, they increase the likelihood of a data breach class action lawsuit in the process. Companies that under-notify may reduce the likelihood of a data breach class action, but could instead find themselves in trouble with government regulators." Wow. What a mess. "For school districts and other entities that suffer data breaches, legal fees and settlements are often among their largest expenses." Yeah. That's a shock.

"Law firms like McDonald Hopkins that manage thousands of cyberattacks every year are particularly interested in privilege," said Schwarcz, the University of Minnesota law professor, who wonders whether lawyers are necessarily best positioned to handle complex digital attacks. In his 2023 Harvard Journal report, Schwarcz writes that the promise of confidentiality is breach coaches' chief offering. The report argues that by inflating the importance of attorney-client privilege, lawyers are able to retain their primacy in the ever-growing and lucrative cyber incident response sector. Similarly, he said, lawyers' emphasis on reducing payouts to parents who sue overstates schools' actual exposure and is another way to promote themselves as providing a tremendous amount of value by limiting the risk of liability by providing a shield.

Their efforts to lock down information and avoid paper trails, he wrote, "ultimately undermine the long-term cybersecurity of their clients and society more broadly." School cyberattacks have led to the widespread release of records that heighten the risk of identity theft for students and staff and trigger data breach notification laws that typically center on preventing fraud. Yet files obtained by The 74 show school cyberattacks carry particularly devastating consequences for the nation's most vulnerable youth. Records about sexual abuse, domestic violence, and other traumatic childhood experiences are found to be at the center of leaks. And hackers have leveraged these files, in particular, to coerce payments.

In Somerset, Massachusetts, a hacker using an encrypted email service extorted school officials with details of past sexual misconduct allegations during a school "show choir" event. The accusations were investigated by local police and no charges were filed. The hacker threatened school officials in records obtained by The 74 by writing: "I am somewhat shocked with the contents of the files because the first file I chose at random is about a predatory pedophilia incident described by young girls in one of your schools. This is very troubling even for us. I hope you've investigated this incident and reported it to the authorities because that is some messed-up stuff." And he didn't say "stuff." "If the other files are as good, we regret not setting a higher price."

Danielle Citron, a University of Virginia law professor, argues that a lack of legal protections around intimate data leaves victims open to further exploitation. She notes that the exposure of intimate records presents a situation where vulnerable kids are being disadvantaged again by weak data security. And of course keeping all of this secret and in the dark doesn't improve data security. Danielle said: "It's not just that you have a leak of information, but the leak then leads to online abuse and torment."

Meanwhile in Minneapolis, an educator reported that someone withdrew more than \$26,000 from their bank account after the district got hacked. In Glendale, California more than 230 educators were required to verify their identity with the IRS after someone filed their taxes fraudulently. In Albuquerque, where school officials said they prevented hackers from acquiring students' personal information, a parent reported being

contacted by the hackers, who placed a "strange call demanding money for ransoming their child."

Nationwide, 135 state laws are devoted to student privacy. Yet they are all unfunded mandates with no enforcement. All 50 states have laws that require businesses and government entities to notify victims when their personal information has been compromised. But the rules vary widely, including definitions of what constitutes a breach, the types of records that are covered, the speed at which consumers must be informed, and the degree to which the information is shared with the general public.

It's a regulatory environment that breach coach Anthony Hendricks, with the Oklahoma City law firm Crowe & Dunlevy, calls "the multiverse of madness." Hendricks said: "It's like you're living in different privacy realities based on the state you live in." He said federal cybersecurity rules could provide a level playing field for data breach victims who have fewer protections because they live in a certain state. By 2026, proposed federal rules could require schools with more than 1,000 students to report cyberattacks to CISA. But questions remain about what might happen to the rules under the new Trump administration and whether they would come with any accountability for school districts or any mechanism to share those reports with the public.

Corporations that are accused of misleading investors about the extent of cyberattacks and data breaches can face Securities and Exchange Commission scrutiny, yet such accountability measures are missing from public schools. The Family Educational Rights and Privacy Act, the federal student privacy law, prohibits schools from disclosing student records, but does not require disclosure when outside forces cause those records to be exposed. Schools having a policy or practice of routinely students' records in violation of FERPA - that's the Family Education Rights and Privacy Act - can theoretically lose their federal funding, but no such sanctions have ever been imposed since the law was enacted in 1974.

The patchwork of data breach notifications are often the only mechanism alerting victims that their information is out there; but with the explosion of cyberattacks across all aspects of modern life, they've grown so common that some see them as little more than junk mail. Schwarcz, the Minnesota law professor, is also a Minneapolis Public Schools parent. He told The 74 he got the district's September 2023 breach note in the mail, but he "didn't even read it." The vague notices, he said, are mostly worthless. It may be enforcement against districts' misleading practices that ultimately forces school systems to act with more transparency, said Attai, a data privacy consultant. She urges educators to "communicate very carefully, very deliberately, and very accurately" the known facts of cyberattacks and data breaches. Okay. So this is all a big mess.

Leo: Yeah, no kidding.

Steve: When an enterprise's security is breached, and its proprietary data are leaked, details of its internal operations, employees, and customers, as we know, can become public.

Leo: I think it has to; right? I think the law requires it; does it not?

Steve: Well, yes. The SEC absolutely requires it. And, you know, heads will roll among those on the board if that doesn't happen.

Leo: Sure.

Steve: There isn't the same thing within our educational system. When personal and private records being kept by U.S. public schools are leaked, as now happens with distressing regularity, disclosure of the private and potentially damaging details of our nation's children hangs in the balance. Administrators of these public institutions fear reprisals from the parents of the students that have been placed in their charge, and also fear the loss of trust that accompanies any acknowledgement of wrongdoing. So expensive specialist law firms and attorneys are now being brought in under the cover of darkness as a means of abusing the attorney-client privilege privacy shield protections. And responsibility is handed over to these attorneys, who are only too happy to take the reins in return for their fat attorney fees.

At this point the school administrators are able to answer any question with "You'll need to speak with our attorneys since they're conducting an ongoing investigation." Which, as we saw, can stretch out for more than a year because, well, you know, these things take time. You can't rush these things. We wouldn't want to over-report or under-report.

Meanwhile, insurance companies are working to determine how to best profit from the panic and the threat of ransomware which has been ignited throughout the public school system. On the one hand, they want to write policies and collect their quarterly insurance premiums. And on the other hand they want to minimize and limit their exposure. The ransomware extortionists are able to use the threat of student body private information disclosure to induce the insurers of these school systems to cough up juicy ransom payments.

So it's always useful when we're able to examine the facts and find some way to see that things will somehow get better. But I'm at a loss here, as I said at the top. Ultimately, taxpayer money is being funneled into the wallets of cybercriminals from insurance companies by way of our nation's public school systems. And I can't see, you know, any functional mechanism for holding anyone accountable. So why would we expect any of this to change?

Leo: Well, I think you can. You do the same thing the SEC does with public corporations. You do it with schools, with public schools, anyway. You can't do it with private schools, probably.

Steve: Heads roll?

Leo: Yeah, you pass a law. This is a data breach. And the subjects of the data breach have the right to know that their information's been compromised. So you're required to disclose.

Steve: It sounds like next year that there will be some federal legislation that may pass.

Leo: You just need expansive data breach legislation that says any time there's a data breach, you have two weeks to reveal it to the people who were the subject of the breach.

Steve: And you can't leave it up to the states because, as these guys said, it is an absolute disaster patchwork. It's just a quilt of overlapping and contradictory regulations.

Leo: Yeah, yeah. But I think you could have a comprehensive federal data breach law. Absolutely. And that's what you need.

Steve: And we don't yet, no.

Copyright (c) 2014 by Steve Gibson and Leo Laporte. SOME RIGHTS RESERVED

This work is licensed for the good of the Internet Community under the Creative Commons License v2.5. See the following Web page for details:
<http://creativecommons.org/licenses/by-nc-sa/2.5/>