



Exposed Cloud Databases

Description: This week we track further occurrences of ransomware in Florida and elsewhere. We check in on the state of the "going dark" anti-encryption debate. We look at a stunning new BlueKeep proof-of-concept demo produced by the guys at SophosLabs. We update some miscellany and present some closing-the-loop feedback from our terrific listeners. Then we examine the nature of the continuing problem of massive publicly exposed databases. In the third example of this just this week, we discover a prolific Chinese IoT manufacturer who is logging more than a million of their customers' devices into an exposed database of two-billion-plus records - which returns us to the dilemma we have with the utter lack of oversight and control over our own IoT devices, and the need to soberly reconsider what "IoT" stands for.

High quality (64 kbps) mp3 audio file URL: <http://media.GRC.com/sn/SN-721.mp3>

Quarter size (16 kbps) mp3 audio file URL: <http://media.GRC.com/sn/sn-721-lq.mp3>

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Yes, yet another ransomware attack in Florida and now Georgia, as well, and what your city should be doing to protect your data. Also coming up we'll take a look at cloud breaches. There's a couple of new ones, including one that you just won't believe. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 721, recorded Tuesday, July 2nd, 2019: Exposed Cloud Databases.

It's time for Security Now!. We're going to save your bacon today with...

Steve Gibson: First, yes.

Leo: The king of bacon saving, Mr. Steve Gibson is here. Hello, Steve.

Steve: Hey, Leo. Great to be with you again for Episode 721.

Leo: Holy moly.

Steve: Counting down to 999 and the end of existence as...

Leo: Can you stop saying that? You're making me sad. I know it's years away, but it just makes me sad.

Steve: Ah, well. We're going to have resolved all of these problems.

Leo: Oh, that's right. Oh, never mind. No problem.

Steve: So, yeah, there'll be like nothing left to talk about.

Leo: Yeah. It's all fixed now.

Steve: I have mentioned before that I deliberately skip over stories about yet another exposed database in the cloud. But three different things happened in this past week, and one of them is a little bit breathtaking, so much so that we're going to redefine what the initials IoT stand for.

Leo: Oh, boy.

Steve: But first...

Leo: Initialization of terror?

Steve: Oh, that's not bad. But first we're going to track further occurrences, believe it or not, further occurrences of ransomware in Florida and elsewhere. We check in on the state of this "going dark" anti-encryption debate. And I stumbled into something that's more than I could cover with everything else we have to do today, so it's probably our topic for next week, which is the Ghost Protocol. I kid you not, that's its official name, which GCHQ has proposed. We're also going to look at a new stunning BlueKeep proof-of-concept demo which was produced by the guys at SophosLabs. We update a bunch of miscellany, some closing-the-loop feedback from our terrific listeners, and then spend much more time than we normally do on this nature of the continuing problem of massive publicly exposed databases on the Internet.

And as I mentioned, in this third example of the three that occurred just this week, we discover a prolific Chinese IoT manufacturer who is logging more than a million of their customers' devices' activities into a massive, more than two billion record, publicly exposed database. They have not responded to calls to shut it down, to protect it from the public.

Which returns us to the dilemma that we were talking about, I think very usefully, last week - although it's a dilemma without a solution, so that's annoying - which is the utter lack of oversight and control that we have now over our own IoT devices and what it means. Which brings us to the need to soberly rename what IoT is an abbreviation for. So I think an interesting podcast for our deserving listeners, and a very apropos Picture of the Week, as it happens.

Leo: We have a number of guesses, speculative guesses in the chatroom for what IoT might now stand for. I think one of them is accurate, so I won't say a word.

Steve: If somebody cheated and looked at the end of the show notes...

Leo: Oh, yeah, it's in the show notes, isn't it, yeah.

Steve: Yeah, yeah.

Leo: What are you laughing at?

Steve: Well, our Picture of the Week, which is not genius, but it's just fun. We have the boss sitting behind his desk, and he has an employee standing in front of the desk who announces to his boss, "Our devices are now 100% secure." And the boss is very impressed. He says, "How did you do that?" All caps. And the employee says, "I turned them all off."

Leo: The only way, baby.

Steve: So, yeah. So speaking of turning them all off, more ransomware news from Florida. We talked last week...

Leo: It's so amazing. It's so amazing.

Steve: I know. I know. We talked last week about a Riviera...

Leo: Riviera Beach was...

Steve: Riviera Beach, yes, in Florida, right. And that was a relatively small community of 30,000. On June 10th, the computers of Lake City, Florida, population 12,000, ceased to function as malware encrypted their data after they were hit by a powerful ransomware attack that was named "Triple Threat" because it mixes three different methods of compromise. As we have always seen, and this happened again, an employee opened a document they received in email which infected the city's network with the Emotet trojan, which later downloaded the TrickBot trojan. And then later that downloaded the Ryuk, R-Y-U-K, ransomware. And Ryuk is something that we're going to be hearing - or maybe it's Ruk? Anyway...

Leo: I'm think, I'm guessing, it's like the Three Stooges "Nyuk nyuk nyuk." It's ryuk, ryuk, ryuk. R-Y-U-K; right?

Steve: Ryuk, ryuk, ryuk. So the Ryuk ransomware, but it's really nothing to ryuk, ryuk, ryuk about. Now, in this case these guys quickly detected the attack, disconnected the affected systems within minutes of the discovery that something was very wrong. But

that, of course, all this stuff moves at the speed of light. So the malware was still able to infect systems controlling most of the city's landline phones and email systems, and of course lots of other ancillary systems. The landline phones and email, being communications, was very much in everyone's face, which forced employees to revert to pen and paper in an attempt to continue working. Fortunately, in this case, the police and fire and other emergency services remained unaffected since they were on a separate network.

They didn't receive a demand for ransom immediately, and they were wondering whether it might have been because they quickly took the systems offline, thinking that maybe the bad guys didn't know they had been infected. But it turns out that wasn't the case. The request was received for ransom a week later by the Florida League of Cities, which is the insurance provider for the city of Lake City, which is the city that was infected.

Leo: Wait a minute. You're saying that the ransomware guys bypassed the middleman and went straight to the insurance company?

Steve: Yes.

Leo: Oh, my god.

Steve: I know.

Leo: Oh, my.

Steve: Which is a little chilling.

Leo: Well, it's a spear phishing attack; right? These are targeted attacks.

Steve: Yes, yes. So this league, the Florida League of Cities, the insurance carrier, began negotiating with the attacker and agreed to pay 42 bitcoin. At the time - bitcoin has been...

Leo: That's a lot.

Steve: ...jumping around a lot recently. It hit a peak at \$13,000, but it just today dropped below 10. But at the time it was 400,000. And I was interested to hear that the attacker sent the ransom demand to the city's insurance carrier. This suggests more sophistication than I had expected, but I think we need to recalibrate ourselves. Now, get this. The city's information technology director, whose name is Brian Hawkins, said: "Our systems are shut down, but there is no evidence to indicate any sensitive data has been compromised."

Leo: Nyuk, nyuk, nyuk.

Steve: Uh-huh, which would be painting a happy face on the disaster.

Leo: There's no evidence.

Steve: Of course, again, he's the city's information technology director, Brian Hawkins.

Leo: Oh. He's the guy responsible for this, then.

Steve: Yes. We'll be coming back to him in a moment. He said: "All customer service payment data, such as credit card data, is stored offsite by third-party vendors and would not have been accessed by an attack like this on our network." However, whoops. With no proper backup to restore the data, Lake City was left with no option but to pay the ransom. Thus the city's leadership approved the ransom payment last Monday, 42 bitcoins were paid last Tuesday, and the city's IT staff began decrypting files later the same day. So pretty quick action, negotiated with the insurance carrier.

The insurance carrier - in fact, I'm getting ahead of myself because I have here written down in the show notes, Mayor Stephen Witt said that the administration made the decision to pay up after talking with the FBI. And the FBI probably said there's nothing we can do to help you. And so Stephen Witt said okay, thank you, next? And they switched over...

Leo: The FBI's standard recommendation is do not pay. But, you know, in this case I think you have to change your tune if you have no recourse because normally they say don't pay for two reasons. It incents the bad guys. Well, believe me, we don't need to incent them anymore.

Steve: Right.

Leo: And it often doesn't work because they may not have or give you the key. But I've been thinking about this. Something so highly publicized?

Steve: Yes.

Leo: You're going to get the key because the ransomware authors want you to think paying works.

Steve: Yes, exactly.

Leo: It may not if you're just a private individual. But if you're a city, I think you have a better shot at getting it back. Don't you?

Steve: Well, yes. And as we talked about last week, when we looked at two previous instances, the Riviera Beach and one in, I think it was in Atlanta, the cost, the actual cost

of not paying - wait. It wasn't Riviera because they were the people who paid \$600,000. So it was two previous ones...

Leo: Baltimore didn't pay, and they were down for weeks and weeks.

Steve: And the remediation cost was millions, many millions of dollars. So although no one wants to pay, unfortunately, it actually ends up being the least expensive proposition. So in this case, for this city, their deductible was \$10,000. So they paid \$10,000. The insurance carrier paid the balance. And the city's IT director, Brian Hawkins, was fired.

Leo: Thank you. Thank you, you nitwit.

Steve: Yes, exactly. So we have 42 bitcoins this week, 65 the week before. So that's in the neighborhood of \$1.3 million for these two cities. So this creates a dilemma.

Leo: There's a third Florida city, by the way, Biscayne Beach, I think, yeah.

Steve: Yup. I'm just getting to it.

Leo: Oh, good, all right.

Steve: So here's the problem. Attackers have apparently found an interesting niche to attack.

Leo: It's a good niche. It's a good niche.

Steve: Like individuals, smaller government municipalities apparently lack sufficient IT infrastructure to allow them to both protect from and recover from targeted attacks. But unlike individuals, they have much more important data at stake, as well as insurance to take the short-term sting out of paying up. And they've got, arguably, much deeper pockets, as well. So, okay. Here's a test. I want everyone currently listening to this podcast to raise your hands if you think - oh, unless you're driving - if you think that the bad...

Leo: Then raise a finger.

Steve: Okay, a finger, good. Yes, take a finger off the wheel if you think that the bad guys have all the money they want and will now stop pestering deep-pocketed, well-insured, small municipalities with targeted attacks.

Leo: They're going to retire now, yeah.

Steve: Okay. Everybody got your hands up? Anyone? No? That's right. We're going to be seeing more of this. And, oh, guess what? That's already happened. Since then - I was trying to, like, wrap up this story.

Leo: I'm sorry, I ruined it for you.

Steve: And news kept coming in. An even smaller village of 3,000 residents in Florida, Key Biscayne, was just yesterday hit. Officials reported, or I guess it was earlier because we have yet another one, officials reported a Ryuk ransomware infection last week, but they haven't decided yet if they want to pay the ransom demand. A special session of the Village Council of Key Biscayne was held last Friday night, with a second session on Saturday, to determine how to respond to the Ryuk ransomware attack on the village's systems. No answer yet. What's happened is, as a consequence of this now sort of ongoing series of high-profile attacks, is that the public and the media are beginning to turn negative on city officials who fail to secure networks and then decide to pay hackers.

Leo: Yeah. In fact, there's a term now, "Riviera Beached." You're been Riviera Beached. So that'll give you some idea.

Steve: Yeah. So now paying ransom is viewed as a sign of a city administration's failure and weakness, rather than being a quick fix to get access back for the citizens' data and so forth. And as we would expect in following all of this up, forensics has always determined that it was a city employee who ushered the malware into the network by opening a malicious email and clicking on something.

Leo: But you can't blame the employee so much. Why is that malware getting through their intrusion detection systems?

Steve: Right.

Leo: Why is it passing - you know, have they been trained? You can't completely blame the employee. This stuff's hard to spot. By the way, there's an update on this from Ars Technica.

Steve: Okay.

Leo: According to Ars, Riviera Beach, in addition to the \$600,000 they spent, will spend a million dollars in remediation. So this is - they're not done.

Steve: Wow.

Leo: And this morning the Georgia Administrative Office of the Courts confirmed they've been hit by Ryuk.

Steve: Yes. And in the show notes, oh...

Leo: It was in the notes? Okay.

Steve: And just as I was trying to wrap up the story...

Leo: Just as you were doing it. Unbelievable.

Steve: ...the state of Georgia's Judicial Council and Administrative Office of the Courts fell victim to a Ryuk ransomware attack.

Leo: So it makes sense because these smaller cities, they don't have budget.

Steve: Yes.

Leo: They have valuable information.

Steve: Yes.

Leo: Plus I thought about this. It's probably easier to spear phish them because they all have public websites, so does the Georgia court, where you can go and find the name of the city manager, his email. You can find the name of the comptroller, her email.

Steve: Ah, very good point.

Leo: So you don't have to do a lot of snooping around to get everything you need.

Steve: And you could probably generate email, like for example spoof the sender data of some employee to be their boss.

Leo: Right. Right.

Steve: Because the managerial structure is also public knowledge. I mean, it's all public record. So you do some research, and you generate very convincing-looking email which you then send to an unwitting employee and say, you know, "Hi, Betty. Following up on our conversation recently, here's the data that I told you I would provide."

Leo: Yeah. You're the city manager, and you send a note to the comptroller saying this spreadsheet looks like the numbers are funny. Can you check it?

Steve: Uh-huh.

Leo: That's all you have to say.

Steve: Yup.

Leo: And unfortunately, you still read news reports saying watch out for email from people you don't know. Which is exactly the wrong thing to watch out for. It's always from people you know.

Steve: Right. So their infrastructure is publicly available. They are insured, so they've got, I mean, and one wonders how long insurance is going to be making these payments before either the rates start to skyrocket or the insurance companies say hey, you know, we're not going to insure you against ransomware. You need to do training, and you need to put systems in place, and also seriously be able to back up and recover yourself from this kind of problem.

Leo: Yeah.

Steve: I mean, it really is, I mean, so their deep pockets...

Leo: My guess, though, is this is not - the insurance company, in this case, what did you say it was, the Florida state governments, Florida city governments?

Steve: Yeah, so maybe - they called themselves the...

Leo: Association of Florida City Governments, something like that.

Steve: Yeah.

Leo: It sounds like it's the governments themselves.

Steve: Florida League of Cities, you're right. So it sounds like it's a private insurance pool that they all...

Leo: Right, exactly.

Steve: Yes.

Leo: Nevertheless, you don't really - how much is left?

Steve: Yeah.

Leo: Wow. Wow. This is - of course. And you know what, if you're the Petaluma City Council or the Irvine City Council, your fiduciary duty means you get to go ask your IT director, what's the situation? Now, Steve, there's something I wanted to ask you about this. It seems like a simple backup would solve this problem. Is that not the case?

Steve: Yes, except, well...

Leo: It has to be offline because you can't have the virus infect it.

Steve: Correct. And the viruses also look for all connected drives and network drives. So, I mean, so if you could do a full nightly image of everything, then...

Leo: And take it offline.

Steve: Exactly.

Leo: You take it offline, or put it somewhere like Glacier, where it's immutable. If you did that - and I think that that's not so hard or expensive.

Steve: And I think the problem is it's just a matter of doing it. Like, for example, we still have Flash in browsers. Like against all reason. Why? Well, because we've had it before. How long has it taken us to stop using old hashes and old ciphers? And so the problem I think also is none of these municipalities have as much budget. I wouldn't be at all surprised if the IT people are saying to their boards...

Leo: Please, buy us a hard drive. Please, I beg of you.

Steve: Yes. Like we need this much money. And I'm sure the city council says, we don't have it. I mean, yeah, we'd love, oh, we'd love to repave our roads, and we'd like to update our parking meters and blah blah blah, you know. I mean, so I'm sure it's just a matter of, like, the stuff is working. And so, like...

Leo: Yeah. That's really what it is. By the way, a user in our chatroom just sent me a link to the Florida League of Cities Municipal Insurance Trust. On the front page: "Take our webinar on remediating and avoiding ransomware." On the front page. And they're doing a webinar, but it's not till next week. "Preventing Ransomware: Tools, Tips, Techniques, and Technologies." Lots of T's in there. I hope people...

Steve: Again, it's like - so certainly preventing the infection is what you would like to do. That's probably the least expensive thing you could do would be to somehow keep this from happening. But in a large organization, I mean, a large municipality, how do you get that message out to everyone?

Leo: No, it's hard.

Steve: Especially when they've got bad habits from home and from their phone and from everywhere else on the Internet. And it only takes one. That's the other problem is, you know, back when we were talking about Sony and the advanced persistent threat there, I said on the podcast, I don't want that job.

Leo: No.

Steve: You cannot secure something that big. You just can't.

Leo: On Saturday we had, in studio, the man in charge of IT at West Point Military Academy. And I asked him about this because we were talking about this very story. And he said, "Well, we have the U.S. Army Cyber Defense Command on campus, and they're very good." He said, "But the problem is you only have to make one mistake. We have to be perfect. The hackers don't have to be perfect. We have to be perfect. No mistakes." And that's hard to do, no matter what, even for West Point.

Steve: Well, and consider, okay, in the case of the last two attacks, average of half a million dollars ransom.

Leo: For crying out loud. Buy a technology. Put it in the closet.

Steve: That is a serious ransom windfall for the bad guys. And how many cities are there in the United States? I mean...

Leo: Oh, yeah. And they're all this bad, I'm sure.

Steve: Yeah.

Leo: Every one of them.

Steve: Well, because, again, it's aging infrastructure. It's systems that, I mean, they're probably running - some of them are probably running XP still because they just - it's like, oh, this works. Or our city infrastructure was written for XP, and it won't run. It's 16-bit Windows code, and it won't run on Windows 7 or on Windows 10. So they're stuck.

Leo: So, IT guy. Run down to Costco right now, buy a 4TB hard drive, USB. If you have to walk to every machine and copy the data off of that machine and then put that hard drive in the closet, do it now, and do it every week from now on; right?

Steve: At least.

Leo: Seems like.

Steve: Yeah.

Leo: Criminy.

Steve: Well, and I've mentioned that, because it is such a clear issue, and because ransomware will jump to and encrypt anything connected...

Leo: Has to be offline.

Steve: Yes. The system I have built dynamically mounts and images and then unmounts drives which are not visible except during the brief window of time when I'm doing the imaging because...

Leo: And you have to rotate the images in case the most current image copies the encrypted data. You have to be able to go one back, yeah.

Steve: Exactly.

Leo: I don't think it's that hard. I think you've got to do it. It's 88 bucks for a 4TB drive now, \$88.

Steve: Well, and it's just - and it's intent. I mean, it's like, time. Like the IT guy is probably busy keeping his system up, rather than - you know. And then City Council says, okay, you need to tell us that we're completely immune to a ransomware attack.

Leo: Nope. Nope.

Steve: And the IT guy says, "We're not. I've been telling you for two years." You know? And so, I mean, so maybe some money gets freed up. But again, as you said, it takes one mistake, and it takes down a municipality that has insurance and deep pockets. They're going to pay the ransom because it's going to be far more expensive. I mean, they're basically gone until they, you know, like limping along until they pay the ransom. And exactly as you said, in this case the decryption keys were provided same day. Here's your bitcoin. Here's your decryption key.

Leo: The guy's probably down the street; right? I mean, he's probably in the area. Maybe not, but it might as well be. Something else to mention is a good email system like Office 365, Microsoft Exchange, or Gmail will, in most cases, you're not going to get Ryuk through your email. You're just not. But what they can't filter out against is links, infected websites. I mean, they can, but they're not going to be up, you know, I can make today a website they wouldn't know about till tomorrow.

Steve: There is an attack right now. I didn't get it into the show notes. But it's leveraging JavaScript obfuscated in images which are being served by ad servers to unwitting websites.

Leo: How are you going to avoid that? Boy.

Steve: Also just don't click on an ad. Who ever clicked on an ad, anyway? I don't know who clicks on ads. Only bots click on ads.

Leo: Well, couldn't these infect you without your clicking? I mean, they're just running code. Couldn't it just, if you had the flaw - you'd have to have a flaw in your browser; right?

Steve: Yes.

Leo: So remember that Firefox zero-day.

Steve: That's right. You do not want to - because browsers are sandboxed, and so they're trying not to let anything out without some user action. And so, I mean, there's a big difference between viewing something and clicking something. Clicking lets everything loose because you don't know what's behind that button.

Leo: Yeah, yeah. Wow. It's a scary world.

Steve: Meanwhile, the "going dark" problem that we touch on from time to time is not going away. And I don't think it shows any signs of going away. A headline in Forbes caught my eye. The headline read, and just I'll say first that it's not true, it read U.S. - well, okay, no. Technically, it is true.

Leo: It's not - it's a rumor.

Steve: Yes. "U.S. May Outlaw Messaging Encryption Used by WhatsApp, iMessage, and Others."

Leo: See, I've been watching this story with interest, too. And I haven't talked about it on any of our shows because it's just a leak. It's a rumor.

Steve: Yeah. Right.

Leo: But the minute, and I wouldn't put it past President Trump, the minute that they propose this, we've got to raise holy hell.

Steve: So, okay.

Leo: But they haven't yet.

Steve: So Forbes saying, you know, normally pretty upstanding, although it wasn't an author who I know...

Leo: Forbes is not upstanding. That's the problem is half their content's good; half their content's not.

Steve: Well, so this is "U.S. May Outlaw." Well, okay, yeah. It may.

Leo: Yeah, may, yeah.

Steve: Anyway, so I thought, what? How did I miss that? And so of course the headline was more speculative than anything. Anyway, I tracked down the source material for this article, which was a story in Politico, which actually had some useful reporting. The story was titled "Trump Officials Weigh Encryption Crackdown." I think even that is a little inflammatory. But of course, you know, click bait.

So I've trimmed down what Politico wrote, removing a lot of the superfluous stuff and fluff, like of them explaining what end-to-end encryption is that our audience of course all knows. But there were some interesting bits left which are useful for getting a sense of the environment. So paraphrasing heavily:

"Senior Trump administration officials met last Wednesday to discuss whether to seek legislation prohibiting tech companies from using forms of encryption that law enforcement can't break." The encryption challenge - which the government, as we know, calls "going dark" - was the focus of a National Security Council - that's NSC - meeting Wednesday morning, last Wednesday, that included the number two officials, the so-called "deputies," from several key agencies, according to three people who are familiar with the matter.

So Politico says: "Senior officials debated whether to ask Congress to effectively outlaw end-to-end encryption." One of the people interviewed said the two paths were to either put out a statement, a general position on encryption, and say that they would continue to work on a solution; or, second path, to ask Congress for legislation.

The unreported meeting of the NSC's so-called Deputies Committee did not produce a decision. "A decision to press for legislation would have" - this is Politico writing - "far-reaching consequences for the privacy and security, effectively forcing companies such as Apple and Google to water down the security features on their smartphones and other devices. A ban on end-to-end encryption" - which is not really what anyone is proposing - "would make it easier for law enforcement and intelligence agents to access suspects' data; but such a measure would also make it easier for hackers and spies to steal Americans' private data by creating loopholes in encryption" - which is more correct - "that are designed for the government, but accessible to anyone who reverse engineers them." Okay, maybe. "Watering down encryption would also endanger people who rely on scrambled communications to hide from stalkers and abusive ex-spouses," blah blah blah, right. So, you know, this is just...

So Politico says: "Politico was unable to determine what participating agency leaders said during the meeting, but there is a well-known fault line on encryption within the executive branch. The DOJ and the FBI argue that catching criminals and terrorists should be the top priority, even if watering down encryption creates hacking risks. The Commerce and State departments disagree, pointing to the economic, security, and diplomatic consequences of mandating encryption backdoors.

"DHS [the Department of Homeland Security] is internally divided. The Cybersecurity and Infrastructure Security Agency" - the CISA we've talked about - "knows the importance of encrypting sensitive data, especially in critical infrastructure operations, but ICE and the Secret Service regularly run into encryption roadblocks during their investigations. An NSC spokesperson declined to comment on the meeting."

So the high-level NSC discussion highlights "how policymakers have continued grappling with encryption even as it has receded from the headlines." And there's a lobbyist familiar with the discussions was quoted: "There is a significant administration-wide effort underway on what to do about the 'going dark' issue."

And then, finally, they said tech companies such as Apple and Google - okay, we know about this. San Bernardino, we know about that. Oh. "The transition between the Obama and Trump administrations saw a hand-off of sorts between two high-profile advocates of the need to access encrypted data. After President Donald Trump fired Comey, Deputy Attorney General Rod Rosenstein succeeded him as the government's top 'going dark' warrior. Rosenstein, who dealt with the issue as a U.S. attorney, warned vaguely that cooperation with Silicon Valley was unlikely to work, implying that legislation might be needed. But now Rosenstein is gone, and experts generally agree that Congress is unlikely to pass a bill requiring warrant-compatible encryption."

And that's the first time I've seen that term anywhere, and that's interesting because that's sort of what we've been talking about, warrant-compatible encryption, meaning you still have encryption, but it's some means for, under warrant, allowing decryption. Which the reason it's interesting, it's an interesting term, is that we've talked about the Constitution's protections against unwarranted search and seizure that we have in the U.S. But under a warrant you're able to - if you're able to get a judge to issue a warrant, then law enforcement is able to do, with probable cause, to do what they feel they need to. And as we know, senators Richard Burr and Dianne Feinstein floated a draft measure back in 2016, after the San Bernardino issue, but it was met with an intense backlash which quickly killed any prospects for it going forward.

So the presumption is that this was sort of a sentiment-gauging meeting which might be preliminary to more public and open talks. And in any event, as I followed up other links from this story, I discovered the Lawfare Blog, with its back-and-forth discussions about a proposal by the U.K.'s GCHQ.

Leo: Oh, yeah.

Steve: Which I referred to at the top of the show.

Leo: The Ghost Protocol.

Steve: Yes, the Ghost Protocol. We don't have time to talk about it this week. But I'm preliminarily titling next week's podcast "The Ghost Protocol" to discuss it at some length because it's the first instance of a serious government-sourced quasi-compromise beyond where we are now, which is sort of an all or nothing. And there's been immediate pushback from the tech community. I mean, basically the tech community is going to push back against any what they perceive as weakening of encryption. But there's no way around the fact that warrant-compatible encryption is a weakening of encryption.

And so, I mean, it's difficult to see how we move past this. But I think we're going to have to. I mean, either we decide that we need to get shims installed into the devices

before the encryption, as we've talked about, so that we're not weakening the encryption, but the problem is what law enforcement wants is everything to be recorded and then to be able to go back into previously recorded encrypted conversations and retrospectively decrypt them. And so the idea being, well, we don't, you know we don't know ahead of time what the conversations were in, for example, the case of the San Bernardino attack. They wanted to get the guy's smartphone and then look at what conversations had been had prior to the attack. So the fact that those are tantalizingly available is part of the problem.

So I don't know. I mean, this is going to create a tension, and I think this is one of the most interesting things happening today in the security and privacy arena is this problem, this tension between the fact that we now have math that cannot be cracked, and governments are not happy with it.

So our friends at SophosLabs advanced the state of the art in BlueKeep exploits in a stunning proof-of-concept demo video. It's relatively short because there's no audio. It's all video. I don't think it makes sense for us to play it into the podcast. But I have a link to it. And in fact I was going to, a little bit later in the podcast, introduce our listeners to a new service. Okay, good. I'm glad...

Leo: I'll play it while you talk, and that way people watching can watch it; right?

Steve: Well, yes. It is amazing. What this demonstrates is they have two machines side by side in the same system. One is in a virtual machine. And they first demonstrate using Remote Desktop Protocol to connect to the second machine. And it requires a login. You get prompted for a username and password. You can't do anything else. You're just stopped there. So then they close that session down, and then they launch their tool, which is their proof-of-concept which they've created, which does some massaging of memory in order to prepare the exploit and the target machine. Once that's finished, they then close that. Basically that changes the target machine over RDP. They then relaunch an RDP session.

Oh, and the other thing it does is, when you bring up the Remote Desktop Protocol desktop, down in the lower left is an accessibility button for those who need accessibility help. Their exploit changes what that is linked to, to bring up an admin system-level console through remote desktop that gives you full access to the machine. So it's just a neat proof of concept. They are not making it public. And again, this is sort of a mixed blessing because this also further demonstrates what can be done by people who know what they're doing. And we know that, unfortunately, there are probably as many people with mal agendas who know what they're doing as with good agendas.

But I mentioned a link. I've been wanting to bring up a GRC-owned link shortener for some time. Our listeners all know that I've been a long-term user of bit.ly, B-I-T dot L-Y. I've had .sc, I've had grc.sc, as in shortcut, for years. And the problem with using bit.ly is that for some time the pattern of the links I was generating were obvious. It was bit.ly/sn-627 or something. And some rascals went ahead and pre-allocated a bunch of future Security Now! episode numbers and related links, just to be mischievous. And so I thought, okay, I need my own shortener.

Anyway, so this is the debut of GRC's link shortener, grc.sc/bluekeep. And that will take anyone who's interested to this demo video, which is worth taking a look at if you didn't just see it in the video on the podcast, grc.sc/bluekeep. And that just redirects to a link on Vimeo, where the guys at SophosLabs have put this. Still no evidence of a worm. And again, to me it doesn't make any sense.

I think we've seen an Exim worm immediately appeared because the exploit was so time-consuming. I imagine this is being exploited quietly, probably, based on everything we're seeing. We're seeing that what the bad guys want is money. They want bitcoin. And so either you extort people by encrypting their networks in municipalities, or you put bitcoin miners on people's servers that are apparently unattended, and no one is paying any attention to, and you mine bitcoin that way as part of a pool. So I would imagine a lot of these systems that have un-updated - we know that it's more than several million have un-updated RDP servers. They're probably now increasingly running coin mining for someone.

I had a tweet from someone about SQRL. Joshua asked, he said: "I have a question regarding a unique use case of SQRL. I'm active duty military and am frequently using controlled systems that don't allow installation of third-party software, including browser extensions. I also work either on underway ships or in areas where cell phones and other electronic devices are prohibited. Is there any way for SQRL to work for me?" He says: "I tried to sign up for the forum, but my account is in moderation."

And so real quickly, if SQRL succeeds, then part of its success, part of what I will consider success is that it becomes adopted by browsers. All we have to have is Chromium, the Chromium Project, to decide, hey, this makes sense. Immediately then - well, and Mozilla, too. Immediately Chrome and Edge get it, and then Mozilla follows, and we've got browsers covered. So the point being that eventually there's no way that this won't end up being built into browsers. It is a natural thing for that to happen. And I've said in GRC's newsgroups that, having finished a client for Windows, if I knew then what I know now, I would have done this as an extension.

On the other hand, when I began this six years ago, we didn't have a common ecosystem for browser extensions. Everybody was disparate, and nobody was synchronized, and the browsers were competing with each other to a much greater degree. So there wasn't a single solution. And there is something to be said for a single client running in a multi-browser system. That is, if you have a Windows client for SQRL, then you can use it with Chrome and Firefox and IE. And so one installation of your client runs across all the browsers, if it's in the OS. So maybe that'll be the other thing that'll happen is it might be moved into the operating system. That's another place for it to be.

So eventually, I mean, so I can get it where, in a locked-down environment, it's not possible to install things to make this work. And if you can't even put it in an outboard phone, then you're going to have to wait until it's, you know, because it does require some computation. I mean, at the same time, this environment wouldn't allow the use of a password manager. So I don't know what these people do. I guess they have a big book of passwords and somehow enter hopefully hard to memorize passwords manually because, if you can't install a password manager, how do you have a different password for every website?

Leo: I don't know if they go to websites, to be honest with you.

Steve: Yeah. And so in that case...

Leo: Do they surf the 'Net?

Steve: That's a very good point.

Leo: That's why you need SQRL, right, yeah.

Steve: So why do you need SQRL anyway? Right. And just so everyone knows, the SQRL forums are moderated. We've got a big team of moderators who signed up in the beginning days. And I'm so grateful to them all because essentially anybody who joins, just in order to - because we don't want this to be a free-for-all and a spam fest and a mess, which unmoderated forums tend to devolve into. So people have to demonstrate that their intent is sincere. And if anyone starts to become a problem, we'll just remove them because the point is for these forums to be useful for SQRL users. They are not intended to be a place for people to spray their graffiti. So everybody is moderated when they start. And then they come off of moderation after they've proven themselves a little bit.

Oh, and GMT also said: "Searched for SQRL app on iTunes. Found this." And then he said, "sql.chat." And I have no idea what that is. But the only way at the moment, because this has not yet been released to the iTunes store, Jeff is still working on it, is to join the forums, and you'll find links there to get on the beta team for the iOS, SQRL for iOS, and then you're able to use it. And that's what I used as part of my SQRL presentation.

Leo: Somebody in the chatroom has said that there is now a plug-in for SQRL for WordPress. Were you aware of that?

Steve: Yes, there is.

Leo: That's great.

Steve: It's still in its early stages. The guy who did the SQRL for Android...

Leo: Daniel Persson.

Steve: Yes. He saw all of the conversation. And so he said he spent between eight and 10 hours. He knows PHP, and so he created a - it's not feature complete. I haven't looked at it closely. But it is working. And I've also heard that our good friend of the podcast Rasmus Vind, who did the beautiful implementation of SQRL for the forums, apparently another listener of ours has asked him to consider consulting to create a formal WordPress plug-in. So that would be great. I would love to have Rasmus do a commercial-grade version of plug-in for WordPress, which would be great. So, yes, lot's happening. There's work on a native SQRL for Linux. Several people are working on macOS versions. So there's a lot going on.

Leo: Nice.

Steve: And, Leo, "I Am Mother."

Leo: I watched it.

Steve: Yes. And I think we talked about it at the end of the podcast. I don't remember if it was after we stopped recording. But you enjoyed it a lot.

Leo: Yeah.

Steve: Yeah. I went looking for a review when someone posted: "In the most recent Security Now!, Steve mentions a link he was going to send to Leo of a review of 'I Am Mother,' written by a guy who created an account on IMDB solely to post about the film." And this person says: "Did he ever share that URL during the podcast or in the show notes? I would love to read that review, but I haven't been able to spot it by idly browsing through all the reviews of that film on IMDB, nor by browsing or searching transcripts or skimming through the episodes."

So when I went looking for that review in order to respond, I ran across another review that read: "Such an underrated piece of sci-fi." And I have a link to that review in the show notes. And he said: "The reviews I read only gave me the impression that this would be a decent thriller, not ... this." He said: "After watching one generic blockbuster after another, it's always so nice to see small, creative films like 'I Am Mother' are still getting made. I'm really curious about how much this film cost because it looks really impressive, but I also have the feeling that it was a very small budget. This film has solid acting and interesting characters. It touches on complicated ideas with a tightly paced narrative from start to finish, and managed to keep me on the edge of my seat till the very end. If you love to treat yourself to some great science fiction, I highly recommend 'I Am Mother.'" And of course our listeners know I recommend it, too.

Anyway, with the warning that this is a complete spoiler, I mean, and the reason you and I were talking about this particular review, not the one I just read, the one I have a link to, and I cannot read, is that this guy who created an IMDB account just to post this, it's because he was annoyed by people not understanding the movie. I mean, and you really do need to pay attention, especially at the end, where some things are said that are, like, aha. I mean, sort of like the whole thing's been building, and a few last pieces of a puzzle click into place.

Anyway, this was going to be my first use of GRC's link shortener, and this is a good use for it. But again, warning, total spoiler. So for those who have seen it, or you know you will never see it, then grc.sc/iam. [Grc.sc/iam](https://grc.sc/iam). And you and I, Leo, both liked it.

Leo: Yeah. And I'm reading this, and now I realize I half understood what happened at the end.

Steve: Yeah, it was well conceived.

Leo: Yeah, yeah, very - but you have to do some math to really understand what it all meant.

Steve: Right.

Leo: I kind of suspected something, but now I understand completely.

Steve: Yup. So anyway, but again, don't go there if you...

Leo: No, no. Watch it. And don't even watch the trailer. Just watch the movie.

Steve: Yes, yes. The trailer, as I said originally, gave away way too much.

Leo: Yeah.

Steve: I mean, again, not the whole thing. But it's like, just better to see the movie.

Leo: Yeah. You'll enjoy it much more, I think.

Steve: And we've had a lot of tweets from our listeners who have said thank you, thank you, thank you for the recommendation. I might have missed it, if not for you guys talking about it. So it's been popular among our listeners.

Leo: Just to address that first review, which is I wonder how much it cost, I think one of the reasons you see a lot of sci-fi take place inside of space ships or in this case inside of a facility is that's cheap. That's really cheap.

Steve: Yes.

Leo: Because they build one set. They shoot the whole thing. Like "Aliens." Cheap. They shoot it in a set, and they don't have to go anywhere. They don't have to do anything. And I'm going to guess that there was a lot of CGI, a lot of computer graphics in there for - because there's - I can't say anything without giving away anything.

Steve: Yeah.

Leo: But there's enough detail that you would think, wow, this must have taken a lot of work. So a lot of times when you make a movie you can either do it with practical effects, physical things, or you do it with computer graphics. And of course sometimes it's cheaper to do it with practicals. But the kinds of graphics they had I would suspect it was CGI.

Steve: Yeah.

Leo: So CGI and a set and two actors? Come on. Cheap.

Steve: Yeah, yeah. And I think there was a lot of running down the same corridor over and over and over.

Leo: Yeah. You noticed that corridor didn't change much.

Steve: Yeah.

Leo: Yeah. You noticed that, did you? Yeah.

Steve: Yeah.

Leo: But that's fine. So what? It was a great movie. A great story.

Steve: It worked.

Leo: Put the money...

Steve: And again, isn't that what a movie is supposed to be? We've sort of gone into this whole computer graphics overload, which is where it's more like a videogame than it is a story with content.

Leo: Yeah. This doesn't appear to be that, at all.

Steve: No.

Leo: No. It's quite well done, I thought. And provocative. As with all great science fiction, the ideas are what make it interesting.

Steve: Yeah. I got a note from, I guess it must have been email: "Howdy Steve and team." He says: "My name is Bob Johnson."

Leo: My name is Team. As Team, I appreciate that.

Steve: I should have made it capital T. He says: "I've been using SpinRite since v5. During that time I was in the United States Air Force and was my fighter squadron's SCM," and he says, parens, "(small computer manager)." He says: "Basically, I was THE IT department. We had a copy of SpinRite v5," he says, parens, "(I cannot vouch for its license), but we used it to recover many a crashed drive."

He says: "So, when SpinRite 6 came out in 2004, I immediately bought one for myself. It was the best money I have ever spent, hands down. So now I'm fervently awaiting v6.1. I have caught you from time to time mentioning your group of testers. I would like to be one of them. Am I correct in my understanding that only those that find their way to your forum/usefeed/hideaway of testers get to beta test? Okay, then. Challenge accepted. When I have time to poke around, I'll find you guys. In the meantime, reserve me a copy of whatever the latest version is. Take care, see you around." He says: "Bob Johnson, avid SpinRite user, Ketonian, and defeater of diabetes, the latter thanks to you and your

selfless research with your own n=1" - meaning a sample size of one - "ketogenic experiment."

Leo: Nice.

Steve: So, very, very cool. And just for Bob and everybody else, I have two places now, for example, where there's SQRL discussion. There's the public web forums, and then there's our old-school, text-only, boring, dry, Usenet newsgroups. I am so much happier there. With a good news reader, which I have, Gravity for Windows, I can mark threads unread. I can leave posts marked. It's sort of built-in project management and dialogue and interaction and feedback in a community. That's where the work of SpinRite 6.1 will be happening. I just can't - I don't have the same sense of presence or management ability in a web forum setting. It just isn't there. And it's not what it was meant for.

So the way this will work is that I'll explain, for people who are interested, how to go over, how to get into the newsgroups. The good news is most people won't want to. And frankly, we have a good set of people ready there already. But anybody with an existing SpinRite license will be able to use that to download whatever it is that is the current version under test, which I'm sure won't hurt anybody's hard drive. We never have. So I don't think that'll happen. But the point is you don't have to participate in feedback in order to be able to play with what I've got. And it is my intention, as I have said, to be immediately producing useful code. So I'll be talking about it, obviously, on the podcast, since there's so much interest in that here. And that's how we'll manage that going forward. But not in the web forums. That just does not work for me. It'll be in the boring, dry, textual newsgroups because it's sort of a built-in project management environment.

Someone calling himself Alan Turing, obviously not, tweeting as @CyberWarrior010, said: "Steve, not sure you have heard about this yet, so I thought I would drop it in your ear. I have found an Exim worm that is, get this, patching the Exim vulnerability. Fun stuff. I have screenshots if you're interested." I didn't follow up but I just thought I would share the fact that, I mean, that's very cool. I mean, it's illegal to do any kind of a worm; but thank you, whoever you are, for breaking the law and patching people's email servers for them. That seems like a good thing, even though don't do that.

James Alseth said - oh. He said, quote, "'There are no bugs.' That's a hell of a claim. What test suites do you run? What are your testing processes?" And I got a kick out of that. Of course he's referring to my bold claim last week that SpinRite has no bugs because I did release 6 in '04, as Bob mentioned, 15 years ago, and I haven't touched it because there's nothing to fix. And I liked his question because it sort of reflects old school versus new school. My testing processes are 15 years' worth of users using it, and tech support standing by to tell me that I need anything to fix. And so it's actual use for 15 years on countless systems and hard drives that have never revealed that there's anything I can do except to advance it to its next version, which of course is my intention.

And lastly, Spencer Dailey. He said: "Hello again, Steve. This is a neat thing about uBlock origin, probably worth mentioning. I remember back when you were deciding to switch from NoScript to uBlock Origin, and at the time it was presented as a monolithic 'Do I run JavaScript or not?' dilemma. Since then, uBlock Origin has added the ability to disable JavaScript on a per-site basis." Which of course is what NoScript used to have. He says: "It's the icon that looks like a script in the bottom right of the popup. It would probably be good," he writes, "to let listeners know that they can selectively disable JavaScript on a site-by-site basis for a pretty big security and performance gain. Cheers, and a shout-out would be cool. Spencer Dailey from Austin."

So Spencer, there's your shout-out and a thanks for making sure our listeners knew that just scripting can be enabled or disabled per site for anybody who is interested in doing that. It doesn't default normally because, again, we've sort of said, okay, you just can't run the Internet today without scripting.

So as we know, I recently commented that I often don't bother to mention the more or less continuous flow of reports of publicly and inadvertently exposed cloud-based databases. These reports are not super interesting. I mean, it's like, yeah, millions of this and millions of that. But there's no real takeaway from each of those. If we had access to the internal details of each breach, so that we could see and appreciate exactly how each of those happened, then that would make for some fascinating and probably teachable moments for the podcast. But normally there is no access to that. All we hear is somebody discovered this big database that was exposed.

But one of the things we do know is that every one of these breaches must come down to two pivotal mistakes. First, probably without exception, the data that, well, first of all, the data was never intended to be made public. Given the nature of the data which is disclosed, it's clearly not the intention of the people who are curating this data that it be publicly exposed. So that says that the database server should never have been bound to the public-facing server interface.

So the developers of the database systems apparently wrongly assume that network-savvy engineers will be setting these systems up and will know better. So they'll bind the server to the local interface, not globally, not 0.0.0.0, which typically means all networks connected to an interface. Rather they'll bind it to the internal LAN side so that it just doesn't answer any calls from the public side.

And then, second, also without exception, we know that all access to the data, even internally, should be protected by some form of strong authentication mechanism, and preferably several forms. Only if both of these absolutely must-do measures have been ignored could we have the problems that we keep seeing occurring constantly out on the public Internet. The server is on a public-facing interface. It's bound to a public-facing interface. And there's no authentication.

So given that being the case, it's really not accurate to call these data breaches, although the tech press often does, I guess because that sounds sexier. It's not a data breach when a web browser is used to access a website because a web server is publishing the website's contents. We don't call that a data breach. So neither is it a data breach when a database has been published by a database server onto the public Internet so that anyone who queries the server will obtain the database's data. It might have been an inadvertent publication of that data, but no breaching of anything was required to obtain access. It was being openly served and publicly exposed.

So this week, once again, several recent data exposures. And it was them all hitting, especially this third one, that I thought, okay, let's just stop for a minute and talk about this problem. So in the first of these three instances, we have more than five million records - five million - containing personal information and medical insurance data exposed in a database belonging to the insurance marketing website MedicareSupplement.com. MedicareSupplement.com is a U.S.-based marketing site which allows visitors to find supplemental medical insurance available in their area.

Last Thursday, researchers said that they found this publicly available. This happened to be a MongoDB database that had been online for several days and was unprotected by any password or other authentication. So again, like this five million-plus-record database of Medicare supplement insurance interested people. Bob Diachenko, who discovered the database along with researchers at Comparitech, told Threatpost, who did some - a lot of people covered this, but Threatpost said, or Bob said to Threatpost:

"While I have not seen any ransom notes or suspicious activity, the IP in question has first been indexed by a public search engine BinaryEdge on May 10. With a MongoDB exposed in the wild for such a long time," he said, "there is a very high likelihood that it has been accessed by others."

Those five million records contained personal and personally identifiable data, including first and last names, physical addresses, IP addresses, email addresses, dates of birth, and gender. So if nothing else, here's five million records of strongly deanonymizing data. And you can imagine, I mean, if we're in a world now where there are parties interested in exactly this, in deanonymizing people for getting where is this email address physically located? What is the first and last name of the person with this email address, their dates of birth and their gender? Well, here's five million instances out on the 'Net waiting.

Also, within that subset, a 239,000-record subset of that more than five million also indicated the areas of interest for the customers, I guess beyond just medical insurance. For example, specifically, insurance for cancer, as well as life, auto, and supplemental insurance. So there was additional data. Sounds like it was a big marketing database maybe that the MedicareSupplement.com was using to say, hey, folks, come over and take a look at our stuff. Anyway, the researchers disclosed the presence of this vulnerable database to MedicareSupplement.com, after which the database access was disabled and then proper security was put in place. So there's, again, this is an example of the stories that I'm constantly stumbling over and thinking, okay, we don't have - this is only a two-hour podcast. We don't have time.

Leo: Only.

Steve: To talk about, like, this constant problem with databases in the public. Second instance. Also in the news this past week were three, well, and this generated a lot of headlines because Netflix, TD Bank, and another one I have in the show notes I'll get to. But some big names were in this. But this is huge. Okay. So three publicly accessible cloud storage AWS buckets, apparently under the control of an Israeli-based data management company, Attunity, A-T-T-U-N-I-T-Y, leaked, get this, more than, and we don't know how much more than, but more than a terabyte, yes, that's more than a thousand gigabytes of data belonging to half of the Fortune 100. The Fortune 100 are the hundred largest companies in the U.S.

So more than half a terabyte, I mean, more than a terabyte - and it's just that a terabyte was sampled - with half of the Fortune 100 and apparently about several thousand other less than the top 100 companies, publicly leaked to the Internet, including internal business documents, system passwords, and sensitive employee information. UpGuard was the group who discovered this open data, wrote up their discovery under the title "Data Warehouse: How a Vendor for Half the Fortune 100 Exposed a Terabyte of Backups."

They wrote, and I've paraphrased this a bit and cut out a bunch of the stuff that didn't make sense for the podcast: "The UpGuard Data Breach Research team" - and again, not really a data breach when you don't have to breach anything in order to get the data - "can now disclose that a set of cloud storage buckets utilized by data management company Attunity have been secured from any future malicious action." Okay. We don't really know that, any current malicious action. Maybe they'll go open again, who knows. "Attunity, recently acquired by business intelligence platform Qlik" - and I guess that's how you'd say it, it's spelled Q-L-I-K, so probably they're trying to be, you know, Qlik, Q-L-I-K - "provides solutions for data integration." Yeah. "An UpGuard researcher discovered three publicly accessible Amazon S3 buckets related to Attunity."

And I'll just stop here because - to editorialize for a minute about Amazon S3 buckets. You have to work to publicly expose them. I use Amazon. I use Amazon S3. And, for example, I'm hosting videos, SQRL how-to videos, on the SQRL Forums through Amazon AWS. And it's work to make these things public. They're not public by default. So I'm always curious when I see, oh, yeah, there was an Amazon bucket exposed. Well, someone did that on purpose. Maybe they pushed the wrong button? I don't know. But, I mean, you have to try. I just wanted to say that. It's not like it happens by chance.

Anyway, so they said: "Of those three, one contained a large collection of internal business documents." And oh, my god, Leo, I thought I had the link in the show notes, and I don't, because you would be scrolling through, and your head would be spinning if you saw what was there. I mean, like, oh. Anyway, they said: "The total size is uncertain, but the researcher downloaded a sample of about a terabyte in size, including 750 gigabytes of compressed email backups." Think about that, 750GB of compressed email. Email compresses down to nothing in sizes because it's typically text. So 750GB of email, that's like everything from the dawn of time for, like, a huge number of people, 750GB.

They wrote: "Backups of employees' OneDrive accounts were also present and spanned the wide range of information that employees need to perform their jobs: email correspondence, system passwords, sales and marketing contact information, project specifications, and more." In other words, this was an archive of what we were just talking about, Leo, for being protected against ransomware. These were the encrypted, well, I'm sorry, not encrypted. They were the backups of people's systems that were exposed in this Attunity bucket.

"On May 13th, 2019," they write, "an UpGuard researcher discovered publicly accessible Amazon S3 buckets named 'attunity-it,' 'attunity-patch,' and 'attunity-support.' The oldest files in 'attunity-it,' where the bulk of the sensitive data was stored, were uploaded in September of 2014" - so nearly five years ago - "though that does not necessarily mean that they were publicly available at that time. The most recent files had been uploaded just days prior to the researcher discovering the bucket." So the bucket was in active use at the time of its discovery and contained at least five years of archived customer data, meaning the data was current and historical going five years back.

"The researcher," they write, "notified Attunity on May 16, 2019. As a result of time zone complications," they said, "and due to Attunity having been recently acquired by Qlik, the researcher ultimately wound up speaking on the phone with Qlik support. By the next day, public access to the buckets had been removed.

"In previous cases," they wrote, "we have discussed the complexity" - I love this. "We've discussed the complexity of fully analyzing and describing data stores that contain copies of users' workspaces and mailboxes. Work lives are spread across countless files, each one with some importance, but difficult to reduce to any single metric of significance. While various tools can help search across large data stores, gaining those quantitative measures comes at the expense of analyzing the qualitative nature of the data." In other words, there was too much data. Sorry, but 750GB, like we didn't know what we had because there was too much to look at.

They said: "There may be a large number of email addresses, for example, but whether they are enterprise customers or merely marketing targets changes the significance of their inclusion. A few key examples from the Attunity data set can illustrate the kinds of data that users have access to and which can be exposed by misconfigured storage of those users' file collections, as in this exposure of Attunity data." And then they said: "According to Attunity's website, more than 2,000 enterprises and half the Fortune 100 use Attunity."

And, sure enough, a file with a client list found in the repository included a client list with a number of companies commensurate to that description. So they also found, apparently, Attunity's customer list and, oh, look, there's half the Fortune 100 and a couple thousand others. They said: "Every business must have customers, and having commercial relationships entails some exchange of information, leading many data exposures to involve third-party data for customers or other involved entities.

In this case, Attunity's business in cloud migration and data integration also involves supplying and managing the software that processes customer data. Exhaustively documenting the files associated with each of thousands of companies is not feasible or necessary for the research team's purpose of raising awareness of the risks of data leaks. As in past reports, a few examples can highlight the kinds of information that can be exposed through misconfigured storage, and were exposed in the case of Attunity."

Okay. So their report then shows page after page of horrifyingly exposed data, even after having been heavily redacted for their own responsible disclosure, things like password reset emails, mail saying "This is your VPN PIN and password," and much, much more. The idea of all of this data having been publicly exposed for who knows how long is sobering. This is the internal data of half the Fortune 100 and thousands of other large companies who were using this service. And it was inadvertently public.

And of course, depending upon whose hands this fell into - and we don't know because it was there, and UpGuard found it; and then, after they reported it, whoops, it disappeared. But we don't know how long it was open. We don't know who and how many people sucked out these terabytes of internal company data exposing thousands of companies' archived emails, OneDrive backups, and workstation contents. We just don't know. The tech press noted that among these companies were Netflix, TD Bank, and Ford.

Anyway, in conclusion, UpGuard's write-up said, as they're concluding, they said: "Attunity's business is to replicate and migrate data into data lakes for centralized analytics. The risks to Attunity posed by exposed credentials, information, and communications then are risks to the security of the data they process. While many of the files are years old, the bucket was still in use at the time detected and reported by UpGuard, with the most recent files having been modified within days of discovery.

"The chain of events leading to the exposure of the data provides a useful lesson in the ecology of a data leak scenario. Users' workstations may be secured against attackers' break-in, but other IT processes can copy and expose the same data valued by attackers. When such backups are exposed, they contain a variety of data from system credentials to personally identifiable information. Data is not safe if misconfigurations and process errors expose that data to the public Internet."

And I think that was a great point. The point here is that attackers needn't try to seduce their way into an enterprise's workstations to look around when their full backups are available for the taking due to the misconfiguration of an affiliated third-party cloud services provider. I mean, this was horrible. And we have no idea how long it was there before UpGuard found it and said, "Guys, your backdoor is way open." Yikes. And we're not done yet.

Leo: There's more? But wait.

Steve: Yes, but wait. vpnMentor reports "Orvibo Smart Home Devices Leak Billions of User Records." They wrote: "Our expert cybersecurity research team, led by Noam Rotem and Ran Locar, discovered an open database" - once again, on the Internet, open

database - "linked to Orvibo Smart Home products." And Leo, go to Amazon and put in "Orvibo Smart Home" into the Amazon search. O-R-V-I-B-O. These are real people. The database, get this...

Leo: Despite the name.

Steve: Yes, despite the name, they're apparently well known. So this open database includes two billion logs that record everything from usernames, email addresses, and passwords, to precise geolocations.

Leo: Oh, man.

Steve: As long as the database remains open, the amount of data available continues to increase each day. "Orvibo claims to have around a million users. These include private individuals who connected their homes, as well as hotels and other businesses with Orvibo Smart Home devices. This constitutes a massive breach of privacy with far-reaching implications," these guys write. "The data breach affects users from around the world." They said: "We found logs for users in China, Japan, Thailand, the U.S., the U.K., Mexico, France, Australia, and Brazil." They wrote: "We expect that there are more users represented in the two billion-plus logs.

"We first contacted Orvibo via email on June 16. When we did not receive a response after several days, we also tweeted the company to alert them to the breach. They still have not responded, nor has the breach been closed." So in this case we have a company producing IoT devices, smart home devices, with more than a million active users, and a log which is daily being posted to by these devices numbering more than two billion records at this point.

They said: "The amount of data available from Orvibo's servers is enormous. It's also highly specific, which shows just how much data smart home devices can collect about their users. According to the company, there are over a million users who have installed Orvibo products in their homes and businesses. The Chinese company, based in Shenzhen, manufactures 100 different smart home or smart automation products. Data included in the breach" - and Leo, in the show notes I have some redacted screenshots from the logs - "email addresses, passwords, account reset codes, precise location, IP address, username, userID, family name, family ID, smart device, device that accessed the account, scheduling information." And so I have a picture that's been redacted showing this information. And without the redaction, there's the actual stuff.

And they say: "In this first example, we can see that Orvibo is collecting a large amount of data about its users. In this case, not all of the data points are recorded; however, we have other examples that include very specific geo-data, chosen family names, usernames, passwords, and the reset codes that would allow for account takeover." And they show another screenshot of that, and then a screenshot that follows.

They say: "These data logs are for the same account, which we can verify with the matching email address and user ID number. In the first, we only have the email address, IP address, and a reset code. With this code accessible in the data, you could easily lock a user out of their account, since you don't need access to their email to reset the password. The code is available for those who want to reset either their email address or password. This means a bad actor could permanently lock out a user from their account by changing first the password, then the email address. Orvibo does make some effort to concealing the passwords, which are hashed using MD5 without salt.

"The above example is a small sample" - okay, remember, more than two billion of these records - "a small sample of the kind of geolocation data we have. Orvibo keeps logs of precise longitude and latitude coordinates, spelled L-A-T-O-T-I-D-E in the data. The precision of the coordinates can lead us to a user's exact address. This also demonstrates that their products track location in their own right, rather than determining location based on an IP address.

"In this entry from a user in Mexico, it shows exactly which device the user was connected to when the data was logged. According to the Orvibo website, HomeMate is a full smart home system that employs a full range of their products to connect your entire home. This amount of data shows just how vulnerable a user can be should a hacker take advantage of this breach." Data which, again, is right now publicly available.

"One of the products Orvibo offers is a smart mirror. This includes technology to show the weather and display the user's schedule. Here, we have a log for the schedule the user has set with a customized name, 'Winter Week Morning,' giving us precise information about the user's calendar.

"This is a data log that includes a large number of devices connected to a single account. We can see a clear record of the user having one of Orvibo's smart cameras. Another device is named 'massage room.' Though not all of the device names tell us which device is where, it could help someone to pinpoint a device to hack if they wanted to do so. The 'massage room' label also points towards this data likely belonging to a business.

"Another Smart Camera log included a message that was recorded word for word. That opens the possibility of a user revealing even more personal information through their account. It's important to note that not every single data log included every type of personal information. However, even with over two billion records to search through, there was enough information to put together several threads and create a full picture of a user's identity."

Okay. Now, just for a minute, stop to consider that this is a Chinese company, based in Communist Shenzhen China. The concern here is with a massive and incredibly irresponsible disclosure of extremely sensitive customer data. But even if this massive multibillion-record database were not exposed, it is still being gathered and retained. But to what end and business purpose? Why log all of this?

"A breach of this size," they say, "has massive implications. Each device in Orvibo's product catalog can have a different negative impact on its users. This is on top of having an abundance of identifying information about users. Much of the data can be pieced together both to disrupt a person's home while possibly leading to further hacks.

"Though Orvibo does hash its passwords, we tested the security ourselves to see how easy it was to discover the real password. In some cases, we uncovered our own password. In order to test this, we created our own account, then searched for our email address to see what account information was accessible. Though our chosen password was hashed, it was easy to crack." So they set up an Orvibo device, created an account, set up a password. Then they went and pulled the transactions from the publicly exposed transaction database on the Internet, found it, found their hashed password, and reversed the simple MD5 hash in order to verify that that's what was going on.

They said: "If Orvibo had added salt to their hashed passwords, it would have created a more complex string that is far more difficult to crack. Salt" - and blah blah blah. We know about salt. But these are unsalted MD5 passwords.

They say: "There are Orvibo devices whose poor security could have severe consequences. A number of the devices offered by Orvibo fall under the umbrella of

'home security.' They include smart locks, home security cameras, full smart home kits. With the information that has leaked, it's clear that there is nothing secure about these devices. Even having one of these devices installed could undermine, rather than enhance, your physical security.

"There's enough information leaked from the database that it makes taking over a user's account a simple task. A malicious actor could easily access the video feed from one of Orvibo's smart cameras by entering into another user's account with the credentials found in the database. At the same time, it would be easy to unlock a door from the same account. With precise geolocation, this simplifies home break-ins, an event smart homes are supposed to help protect against."

And Leo, of course this perfectly dovetails on what we were talking about last week, how utterly blind we are to everything having to do with today's IoT technology. You know, this is a big deal. A huge number of people are installing interconnected, networked, essentially black boxes all over their homes, pressing a few buttons, and presto, look. You know? I can check the dog while I'm away from home, see what the dog is doing. And of course so can someone who may not have your best interests at heart.

The industry has recently been talking about a worm, as we know, this notion of a BlueKeep worm, with great concern, so much so that, as we've talked about it, Microsoft has issued multiple warnings. The U.S.'s NSA and DHS have done the same. But all of those two-plus billion Orvibo transaction log records were created by machine so they could be read by machine and acted on instantaneously, en masse, by machine. So if someone in Communist China wished, or somebody who grabbed that database off the public Internet wished to, an attack could be launched against a million of Orvibo's customers instantly. It's horrifying.

Anyway, so that brings us to my somewhat serious proposal for a renaming of what IoT stands for. I don't think it should be Internet of Things. I think IoT should stand for Installation of Trojan.

Leo: Yup.

Steve: And again, why are they logging billions of transaction records of their million-plus customers? Why? What is the business purpose of keeping that kind of log? As an IoT vendor, we're trusting them. We press some buttons and, oh, look, all of our stuff is connected to the cloud, to something. Orvibo in Shenzhen. We have no control over anything that happens. I mean, and that's important. We have no control. We don't know what is going on. Certainly these million Orvibo users did not know that all of the information that they put into their apps and management, the schedule which they have given to their Magic Mirror that shows their schedule in the morning, is publicly accessible to anybody on the Internet. It is really the Wild West. And what people are doing is installing trojans into their homes.

Leo: And I wonder how much intentionality there is from Orvibo. Like what are they up to? You know? It doesn't seem like it's an accident.

Steve: Well, the logging cannot be an accident. Somewhere, I mean, they know, they're providing storage space.

Leo: They know they're doing this, yeah.

Steve: Yes. Certainly it being public, one imagines, one has to imagine...

Leo: [Crosstalk] intention.

Steve: But they're not even responding to email or tweets. Like they don't care. Incredible.

Leo: Yeah. And even if your stuff isn't from a Chinese company, it's often made in China. So, you know. Who knows what's going on? We've seen [crosstalk].

Steve: Well, we have all of this whole Huawei going on right now with do we trust, I mean, as I've said, I'm surprised other countries are using Windows.

Leo: Right.

Steve: You know? That's just - that's bizarre.

Leo: Right.

Steve: And it does make sense after that scare we had of the Supermicro servers where the question was raised, is there some weird hardware being embedded in the motherboards of these servers? It looks like that was a bogus report, but it could happen.

Leo: Yeah. Okay, well, if it happens, you'll hear it here first. You can listen to us do this show live, if you want. We do it Wednesdays, I'm sorry, Tuesdays at 1:30 Pacific, 4:30 Eastern, 20:30 UTC. There's an audio and video stream always running of what we're doing live or rebroadcasts, if we're not live at TWiT.tv/live. You could choose your favorite stream there. If you are watching, you should probably be in the chatroom or listening at irc.twit.tv, where all the other folks are watching and listening live.

Steve has on-demand versions of the show, both audio and really nicely written transcripts. They take a few days to post after the show is over. Elaine's going to be typing this evening and working on that. You'll find those at GRC.com. While you're there, check out SpinRite, the world's finest hard drive maintenance and recovery utility; progress on SQRL, the SQRL Forums are there. And there's lots of free stuff, useful stuff like a password generator that is second to none, all at GRC.com, the Gibson Research Corporation.

You'll find Steve on Twitter at @SGgrc. And that's one of the ways to reach him. You can direct message him at @SGgrc. You can also go to GRC.com/feedback and fill out the feedback form there. He, I know, likes to hear from you. You can also get audio and video of this show from our site, TWiT.tv/sn for Security Now!, TWiT.tv/sn.

But as always, with all of our shows, we encourage you to subscribe. That's the best way to not only assure that you'll have a copy ready for your Wednesday morning

commute, but that it helps us because we know you're going to download it. You won't skip it by accident. So subscribe in your favorite podcast application.

Steve, thank you so much. We'll see you next time on Security Now!.

Steve: Thanks, buddy.

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/>