Zerologon++

Description: This week we look back at the just-released Chrome 85. We see that an enterprise's choice of VPN gateway really does make a difference. We drop in for an update on what would have to be called the new ransomware gold rush, and we examine the implications of Ring's latest announcement of their flying spy drone I mean webcam. Then we learn how much Vitamin D Dr. Fauci takes, and invite our podcast listeners to lock down their UserID of choice at GRC's new web forums using a non-public URL. Then we conclude with the required big update to the Zerologon story which we began last week.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-786.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-786-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Lots to talk about, including why it's so important to choose the right VPN, and not one that a lot of businesses are using right at this moment. We'll also talk about the latest in ransomware. We don't normally give you reports on ransomware issues; but, man, it's just gotten out of control. Steve will give you an update. And then we'll talk about this huge Windows flaw, Zerologon. It just seems to be getting worse. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 786, recorded Tuesday, September 29th, 2020: Zero Logon++.

It's time for Security Now!, the show where we cover your privacy and security online with this guy. He's so secure, he's wearing tinfoil underwear: Mr. Steve Gibson.

Steve Gibson: It's hot and crinkly, Leo.

Leo: He crinkles when he walks.

Steve: I do not recommend it.

Leo: I am looking forward to it. We're going to have so much fun next Thursday - this Thursday? Next Thursday. This Thursday.

Steve: Two days, yup.

Leo: We're going to have an event. We're doing it in conjunction with ITProTV. It's called Hangover: What IT Will Look Like After COVID. And Steve's going to be, of course, one of our panelists with Amy Webb. We've got Don Pezet from ITProTV to be part of this. So it's going to be a lot of fun. I'm really looking forward to it. Thursday.

Steve: That Amy looks like she's going to be a real kick.

Leo: Oh, you'll love Amy. She's been on TWiT many times. She's super smart. I think she will be very provocative, get the thinking going.

Steve: So we're Episode 786 for this last podcast of September, the 29th. I named this Zerologon++ because we introduced the problem of Zerologon last week. And since then, OMG.

Leo: Oh, no.

Steve: It has been described as the perfect Windows vulnerability, that is, if you were rating them in terms of how the bad guys liked vulnerabilities. This thing has just exploded. And lots of news about that. So I thought, okay, that's a great anchor topic for the podcast. But we're going to look back first on the release of Chrome 85 that was just happening last week. We see that an enterprise's choice of VPN gateway really makes a difference. We drop in for an update on what would have to be called the "new ransomware gold rush" as a consequence of the acceleration that we're seeing in attacks.

Also we have to just talk about Ring's latest announcement of their flying spy drone, I mean, webcam, because that's, like, really, guys? Then we're going to briefly learn how much Vitamin D Dr. Fauci takes; invite our podcast listeners to lock down their userID of choice at GRC's new web forums using a non-public URL, so everybody can get their IDs before the general public; and then we're going to conclude with, boy, a big update on the Zerologon story. And really, Leo, a fun Picture of the Week. It's not a blank page this time, not at all.

Leo: Let me look at it here. Wait a minute. What's funny about this? Oh.

Steve: I highlighted something from the label up above.

Leo: Oh. Okay. This is good. This is one like a Where's Waldo for you. And now you can explain this Photo of the Week to me.

Steve: Okay. So this is a photo, for those who are not in front of a screen, of a Seagate BarraCuda hard drive, 2TB BarraCuda. Everything looks fine. Unless you inspect down toward the bottom of the label very closely, Leo. You wonder why Seagate misspelled Singapore.

Leo: Uh-oh.

Steve: That seems a little suspicious.

Leo: That's a giveaway. Singapre.

Steve: Or why maybe it's a product of Thalland, T-H-A-L-L-A-N-D, rather than Thailand.

Leo: Those would be the giveaways right there. But they're in fine print.

Steve: And at that point you'd be thinking, oh, not so sure I want to trust my data to this drive of unknown past. It is the case that you can sort of see there was another label on the drive beforehand.

Leo: Yeah, yeah, yeah.

Steve: And this one was put there. But somebody went to the trouble of forging Seagate BarraCuda hard drives, complete with model numbers, serial numbers, barcodes, I mean, it looks great.

Leo: Why didn't they just Xerox the original? It's crazy.

Steve: Singapore has an "o" in it; you know?

Leo: Yeah, "singapre."

Steve: But not on this drive. So it turns out that we had some fun a while ago talking about - we've talked about counterfeit Cisco hardware. There was that counterfeit device where the researchers kind of went off half-cocked and blamed the manufacturer, but it was actually a counterfeit of theirs that had this bad behavior. Theirs actually didn't. And hard drives have a similar problem. I guess there's been some discussion in the forums about that counterfeiting is a growing problem in the electronics industry. And when you look at this thing, I mean, it's got the Torx screws. Somebody took I guess some other, less high-reputation hard drive. Hopefully it has 2TB.

Leo: That's the other thought, yeah, maybe it's not really 2TB. That happens, too.

Steve: Yeah. You definitely want to make sure like all of those two trillion bytes are actually inside there before you start using it. But anyway, I just thought that was - I got a kick out of this with the misspellings on the label. And just sort of thought I would point out that it is worth carefully inspecting the things that you buy online because there's no quality control happening in the channel. And it's sometimes very difficult to return things from third parties. So yikes.

Last week we noted that Chrome was right on the brink of updating to their v85, but we didn't yet have all the details. And since several of the things that were fixed were extremely serious remote code execution vulnerabilities, we still today don't have all the details, and Google won't be releasing them until they have essentially lost their value as attack vectors by virtue of everyone having updated to Chrome 85. These are not zero-days. That's the good news. And Google stated that they are not aware of any of these being used in the wild. Of course that's not the same as knowing that they are not. But that's all we've got.

As we've observed, that updating process often occurs during a rollout which can take some time. And Google's own Chrome releases document page about this last week starts out by saying the stable channel has been updated to 85.0.4183.121 for Windows, Mac, and Linux, so it's the desktop Chrome, clearly, which they said will roll out over the coming days/weeks. So it might be a while. I would urge anybody listening to this just to, you know, you can normally always induce an update by going into the About page. I mean, I have found that often I've caught it off guard. It goes, oh, hold on just a second, and then it spins some wheels on the screen and then invites me to reload Chrome.

But otherwise, and this is the case with Firefox also, they just don't seem, I don't know, like being preemptive about it, probably because they typically feel that this would be more of an inconvenience to the user. How many times have we had Windows suddenly shut down on us and do an update right in the middle of a presentation or before one or something, where it's been inconvenient?

So anyway, some of the things that were fixed were vulnerabilities that could have enabled zero-click remote code execution on a visitor's machine by simply visiting a hostile site that was aware of and had weaponized the vulnerability. We don't know yet, but even maybe a malvertising ad could have done that to people. And this is why I talk about - this is why we have a separate section at the beginning of the podcast, typically, about browser issues, because this is arguably now the main interface that most users have to the Internet, for better and for worse. I mean, it's the thing we stick out there in our name and hope that it turns out okay.

Google is responding quickly and appropriately, and we know they're on top of their game from a security standpoint. But these things are just - they've become operating systems unto themselves. And they're just that complicated. So overall, 10 security flaws were fixed. And I was glad to see that the top two disclosures earned their discoverers $15,000 each, with the third being rewarded $10,000. And I think this bug bounty model, that and relying upon the ethics of those who discover these problems, has become a crucial aspect of today's threat management landscape.

I mentioned ethics because a weaponizable zero-click unknown remote code execution exploit for the world's leading web browser by a large margin could doubtless have been sold to the likes of Zerodium for a far higher price than the 15 or $20,000 that these discoverers got. So props to them for saying, you know, I'm going to make a little bit of money here. I'm going to feel good about myself and be able to sleep at night because I'm not selling this thing to Zerodium for them to hurt people with eventually, I mean, basically resell this to Zerodium's customers for clearly underhanded reasons.

So I placed a link to the release report in the show notes. I tried to drill down to get additional information on these individual issues, just to see if there was anything more interesting and juicy there. But all the pages came up "permission denied," with a banner across the top reading "Chrome Googlers, we're investigating a problem with issue permissions." So no additional information was readily available. So I couldn't get there. I think they were just having some sort of a site security problem of some sort, which I thought itself was interesting. So for now all we can do is make sure that any Chrome desktop we're using is at 85, and just be glad that these problems were found and

eliminated before their use was observed in the wild, and again that good guys were releasing these things responsibly.

Even if they hadn't sold to Zerodium, the other thing we've seen are irresponsible people just tweeting them. And it's like, okay, that's not good either because it takes Google or anybody some time, days in the best case, to get a fix figured out and up to users. And in fact the story that we'll be talking about at the end of the podcast of the Zerologon problem, this first came to light when Secura figured six weeks had been long enough, two Patch Tuesdays, long enough for them to then provide the details, which they were no doubt really anxious to provide. Who can blame them? They discovered this, and it's cool. And so that wasn't enough. Not last week; and, whoa, not as we'll be talking about in the week since.

And so Leo, it was interesting that ExpressVPN was a sponsor because I had in my show notes here, I said: "From time to time through the years we've had various recommendable VPN service providers as sponsors of the TWiT Network." And of course we have one now. I said: "Our argument then has been that the provider you choose makes a difference." And so this little bit of news is an example of just that.

Two researchers who work for a network security platform provider, SAM Seamless Network, introduced their most recent discovery by explaining, they wrote: "Many have written, and probably more will be told about the dramatic change in the way each of us works remotely following the COVID-19 pandemic." And of course, as you said, Leo, we're going to be doing a panel about this in two days. "As security researchers, we've been trying to assess whether the existing security solutions address the new situation. We noticed that many companies are resulting in requiring employees to connect to the office via VPN. Whether it's because of onsite data centers or IP anchoring, many small businesses - less than 20 employees - use a VPN server that is usually also the company's gateway.

"Not too long after we began our research, the name FortiGate was thrown into the air. We instantly grabbed the FortiGate that we kept as a backup in the office and began exploring. Surprisingly, or not, we quickly found that under default configuration the SSL VPN is not as protected as it should be, and is vulnerable to" - another thing you mentioned, Leo - "man-in-the-middle attacks quite easily. The FortiGate SSL VPN client only verifies that the CA [Certificate Authority] was issued by FortiGate or another trusted CA.

"Therefore an attacker can easily" - actually, they said "CA" and they meant cert here. And we'll get to that in a second. "Therefore an attacker can easily present a certificate issued to a different FortiGate router without raising any flags, and implement a man-in-the-middle attack." And they finish their introduction saying: "We searched and found over 200,000 vulnerable businesses in a matter of minutes."

Leo: Whoa.

Steve: I know. It's just, really? Still? Really? What year is this, Leo? Oh. Okay. So the scenario is, as we know, employees have moved home in droves and are now connecting back to their corporate network over in this case the popular FortiGate VPN gateway, perhaps a quite recently deployed FortiGate gateway, if this was suddenly something that had to happen, you know, in March or April of this year. They describe, these researchers, an entirely feasible scenario where a compromised IoT device in the remote user's home is able to use ARP poisoning to intercept the employee's connection back to the corporate VPN gateway and insert itself into the loop.

Now, these guys are IoT security folks, so they take the evil IoT device attack model. In practice, such interception could be performed anywhere along the path from the employee's VPN client to the corporate VPN endpoint. So, for example, a bonanza could be had by a serious competitor, say in corporate espionage mode, arranging to intercept their competitors' bandwidth close to but outside the company. "Close to" so that they would be able to see a lot of the incoming connections to that company's VPN.

So the essence of the flaw that was discovered was that any device or attacker who can arrange to intercept the client's remote connection can relay their authentication to the corporate FortiGate gateway, despite its being protected by a VPN. The FortiGate VPN gateway is an SSL VPN, so it establishes first a standard SSL/TLS connection using the client and server model that we're all quite familiar with on this podcast. To that end, every FortiGate VPN gateway - so the gateway at the enterprise side; right? - has a unique certificate where the certificate's common name, that is, just like the name on the certificate, like the name on my certs are GRC.com. The certificate's common name consists of the model number and serial number of that particular gateway. Since every serial number is unique, every certificate is unique.

The trouble arises, as these guys note, when the FortiGate VPN client used with its default settings does not bother to check whether it's connecting to the correct FortiGate VPN. It doesn't look at the connecting certificate's name at all. It doesn't bother to check that it's connecting to any FortiGate VPN gateway. If the certificate is valid and has been signed by any valid and recognized certificate authority that that user's PC/Mac/Linux machine, whatever it is, trusts, that's good enough for the FortiGate client.

In other words, this would be like having our web browsers not even looking at the name on the server certificates when our browsers connect to remote web servers offering services. All that would be checked was that the certificate itself received from the web server was valid, and it was signed by someone valid. We all know that would be totally nuts, since any sort of attack to intercept or reroute traffic, like just a DNS spoof, would allow an attacker to redirect traffic to their server, spoof the identity of the server with their own certificate, you know, a valid certificate with somebody else's name, and then go from there.

So the FortiGate VPN offers zero protection against any form of man-in-the-middle interception. An attacker, even an IoT light bulb on the home user's LAN, simply answers the client's TLS connection with any valid certificate. The attacker simultaneously turns around and reaches out to the Fortinet VPN gateway at the enterprise and initiates its own fresh VPN connection. Then when the user's, the valid user's client sends its authentication credentials, right, to log into the network, after establishing what it thinks it established, a VPN connection, it's the attacker who receives those login credentials, probably records them for future use, and then forwards them to the connection it has made to the corporate gateway.

The attack then proceeds, with the attacker having access to the decrypted content of everything that moves back and forth through it. And note that the attacker is now also on the corporate LAN, under that user's credentials. So it can get up to all manner of other mischief while the client is online and none the wiser. And in fact even when the client disconnects, it can maintain the connection and continue doing whatever it might want to do.

So the problem is bad, and the fix is, I mean, it's unconscionably trivial. If the client were simply configured to verify the gateway's certificate by name, and require that that certificate be signed by the FortiGate certificate authority and none other, whose root certificate would have been installed, presumably at client setup, then any intercepting man in the middle would be locked out since they could never authenticate to the client with a FortiGate VPN server certificate because they could never make one. They would

never have the private key required to sign such a certificate. Only FortiGate would have that.

Leo: I wonder if this is by design, rather than a flaw.

Steve: Oh, boy.

Leo: Like they want, like the businesses want to intercept or something like that; you know?

Steve: No, well, so I don't see that. The only way to explain - but that's a perfect segue, Leo. The only way to explain this clearly insecure configuration for a VPN would be that all of the concern must have been about the client authenticating themselves to the corporate LAN. That is, well, yeah, bring up a VPN, that's great. But mostly, once that's up, then the client is going to have to log into the corporate LAN in order to prove they are a client of the company. And so that's necessary, but obviously that's not sufficient protection. The clear need to have the gateway also authenticate to the client was apparently somehow missed.

So being responsible lads, the researchers reached out to the Fortinet folks to share what they had found and to get their comments. Believe it or not, FortiGate reportedly responded that they are well aware of it, but are not going to change it. They claim that since the user has the ability to manually replace the certificate that was provided with the system, it's the user's responsibility to make sure the connection is protected. Wow. And the SAM Seamless Network guys noted in their write-up, they said: "Moreover, there is no clear warning by Fortinet to the user that this major security flaw exists when using the default certificate. Instead, a vague message is displayed."

They concluded their posting, saying: "We decided to take the research one step further and check how many Fortinet devices are vulnerable to this type of attack. Using the Shodan.io database, we found approximately 230,000 FortiGate devices that are using the VPN's functionality. Out of those, roughly 88%, that is, over 200,000 businesses are using the default configuration that can be easily breached using any man-in-the-middle method."

So of course one of my favorite aphorisms on this podcast has been "the tyranny of the default." Right? Fortinet's defense that their gateway can be made secure is horrifying when you consider that the security that any purchaser would assume they're obtaining when they turn on a VPN, you know, what else are they buying? And the numbers shown from the Shodan scan that indeed 88% of all the identified installations are using the insecure default, demonstrates that Fortinet is badly letting their customers down. So, wow. I mean, it's just unconscionable. I'm astounded by it. And so of course I wanted to share it with our listeners. It does matter which VPN solution you choose.

Leo: Yeah.

Steve: So we have what BleepingComputer, who has, as we know, been intensely focused upon the ransomware world since it first emerged into the infosecurity scene, they've titled this the "Ransomware Gold Rush," to describe just the past week. Lawrence Abrams has assembled a timeline of events occurring during the seven-day period from September 19th through the 25th. So the 19th was Saturday before last, to the 25th,

which was last Friday. I want to give everyone a sense for what's actually going on here, for example, during a seven-day period in the land of ransomware. On the 19th, Michael Gillespie and PolarToffee found a new ransomware called Egregor that appears to be a Sekhmet spinoff.

Leo: Oh, Sekhmet?

Steve: I know, the old Sekhmet.

Leo: The old Sekhmet's at it again.

Steve: Duck when that thing comes at you. It uses a random extension and drops a ransom note named RECOVER-FILES.txt.~. Now, recall that Michael Gillespie, he's the guy who has been reverse-engineering those poorly written, but still effective on their surface, ransomware strains that permit some hope of non-ransom payment decryption. His pages over at Bleeping Computer now contain decryptors for 23 different kind of low-end ransomware variants. They're low-end because they can be decrypted without needing to receive the key after paying the ransom to the bad guys.

So that's who he is, and he's very involved. He's become sort of a focal point of a lot of this intel. So you'll hear his name a lot as I run through what happened in this one week. As a matter of fact, he also found a different new variant of the LeakThemAll ransomware that appends .montana to encrypted files and drops a ransom note of !HELP!.txt. GrujaRS found a new ransomware that appends the .zhen extension to encrypted files. That was just on Saturday. On Sunday, Michael found a new variant of the STOP ransomware that appends the .kolz (K-O-L-Z) extension to encrypted files. Sunday was a slow day.

Monday the 21st, a new ransomware named ThunderX was discovered. After its analysis, for those fortunate enough to look for it, a free decryptor has been created. It gets added to the list of those that were not written well. The trouble is, of course, many victims won't know that they could look for a decryptor for the ThunderX ransomware. So maybe they'll end up paying ransom they didn't need to.

Also last Monday, in related news, Nathan Wyatt, previously known as the Dark Overlord, pleaded guilty and was sentenced to five years in prison for his extortion threats where he was threatening to publicly release information from his hacking victims unless they agreed to his extortion demands. So hopefully the underworld will see that and think, ooh, maybe that's not such a good idea.

Also that day, last Monday, Michael Gillespie found a new ransomware that appends the .encrypted - not much imagination there - extension and drops a ransom note named SOLVE ENCRYPTED FILES.txt. He also found a new variant of the Matrix ransomware that appends the .JB88 extension and drops a ransom note JB88_README.rtf. Xiaopao found a new Nefilim (N-E-F-I-L-I-M) variant that appends the .TRAPGET extension and drops a ransom note named TRAPGET-INSTRUCTION.txt.

Which brings us to Tuesday of last week. Luxottica, the Italian owner of the Ray-Ban brand, has confirmed that they were the victim of a ransomware attack.

Leo: Good.

Steve: Which has disrupted work, shutting down operations in Italy and China.

Leo: Luxottica is the worst monopolist ever. They bought up every glasses brand there is. Not just Ray-Ban, everything. And they then raised the prices on frames to hundreds of dollars. So sometimes you don't mind.

Steve: Well, I guess maybe they have the money to pay.

Leo: Oh, they do, yes.

Steve: I think they're going to because one imagines maybe their IT staff is not topnotch. Anyway, meanwhile, and this is interesting, a provider of cyber insurance has begun performing their own security scans during their initial underwriting phase. So they're saying, yeah, we'll agree to cover you, but first we're going to check your security ourselves to see if we find anything obviously wrong.

Leo: Of course.

Steve: Since they, yeah, since they've been doing this they're reporting a 65% reduction in subsequent ransomware claims being made against their clients. Yeah. And Michael Gillespie found a new Matrix variant that appends the .FG69 extension and drops a ransom note named FG69_README.rtf. The source for this stuff has to be out in the wild so that script kiddies are just, like, renaming a few things and then launching it themselves. Xiaopao found a new Matrix ransomware variant that appends the .AW46 extension and drops a ransom note named !AW46_INFO!.rtf. GrujaRS found a new ransomware that appends the .CRPTD extension to encrypted files, and 3xp0rt found a ransomware actor selling a complete ransomware kit - ah, there we go - for $2,000.

Which brings us to Wednesday of last week. A leading government technology services provider, Tyler Technologies, has suffered a ransomware attack that has disrupted its operations. We've talked about the problems QNAP NAS devices have been having. Now they've been targeted by the AgeLocker ransomware, which encrypts the device's data and in some cases steals the victim's files, as well. A new ransomware group has been targeting large corporate networks using backdoors of their own design and file-encrypting malware for the initial and final stages of the attack.

Joakim Kennedy found a new ransomware written in Golang that is pretending to be REvil. What's odd about it is that there's no way for its victims to recover their files since there's no contact information provided. So, yeah. Researchers think that perhaps it's just a wiper dressed up as ransomware for some reason.

Which brings us to Thursday the 24th. A new ransomware campaign named Mount Locker, and I have a feeling we'll be hearing more about this, is underway. It exfiltrates its victims' files before encrypting them and is then demanding multimillion-dollar ransoms. I saw one ransom note from Mount Locker in some other coverage for $2 million.

Leo: Wow.

Steve: Someone named S!ri with an exclamation mark for the first "i" found the new Dusk v1.0 ransomware that drops a ransom note named !#!READ-ME!#!.txt. And JAMESWT found a sample of the new Exorcist 2.0 ransomware. Which brings us to and concludes with Friday. Michael Gillespie found a new STOP variant that appends the .copa extension to encrypted files. He also found another new Matrix variant that appends the .DEUS extension and drops a ransom note named DEUS_INFO.rtf.

So, there's just one week in the life of the ransomware world. But it should give everyone a sense for just how totally crazy and out of control this new bitcoin-enabled cyber-extortion has become.

Leo: There's a hospital system, UHS.

Steve: Yes, that's our story after the next one.

Leo: Okay. I don't want to steal your thunder, yeah. But wow. UHS. It's big. Really big.

Steve: Yup. We'll get there in one second. But first we've got - because that's Ryuk. We also have REvil.

Leo: It's spelled REvil; right?

Steve: Right, exactly. Capital R...

Leo: I wonder if they mean it to be Evil.

Steve: Yes. Yes.

Leo: Yes, I believe they do.

Steve: And its actual name is Sodinokibi. You'll recall we've talked about this before.

Leo: Oh, yeah, yeah, I remember that one, yeah.

Steve: Yeah.

Leo: And hackers don't name these, as you point out; right? Security researchers name them, usually.

Steve: Right. And typically it's from something that's found in the file itself. It'll be like, for some reason, like the researchers are trying to - they need to give it a moniker of some sort so that they can talk about it. So they're like looking around. And typically

something will stand out. There'll be something unique about it where they go, okay, let's call this, I mean, who knows what Sodinokibi was. But that's the name.

Leo: So some string that was in the code or something like that.

Steve: Right, right. So these guys are - they're using that ransomware-as-a-service model.

Leo: Oh, Jesus.

Steve: And they, yeah, they recently deposited $1 million in bitcoin into a Russian-speaking hacker forum to demonstrate to potential affiliates that they're soliciting that they mean business.

Leo: Wow.

Steve: Meaning that there is money to be made by joining their affiliate program. Okay. So just to back up a bit, we've covered it because we covered this when this happened. An increasing number of ransomware operations are being conducted as ransomware-as-a-service, where the ransomware developers develop and provide the ransomware and also maintain and run the extortion payment site. And then affiliates are recruited to actually do the hacking into businesses and use the software provided by the developers to conduct the attack. And in a smart move, the ransomware developers typically receive a modest 20 to 30% cut, and the affiliate gets the balance of 70 to 80 percent of whatever ransom payment they're able to negotiate.

So, and we talked about this at the time. At one point this REvil gang that is using the Sodinokibi ransomware had closed their site to new affiliates, stating that they had all they needed, thanks very much, we're good. That changed yesterday when the gang announced that they were once again recruiting new affiliates to distribute their ransomware. Their posting indicates that they're seeking teams of skilled hackers at penetration or experienced individuals.

In their posting they said: One, teams that already have experience and skills in penetration testing, working with msf - I'm sure that's Microsoft - cs, koadic, nas, tape, hyper-v, and analogs of the listed software and devices; or people who have experience, but do not have access to work. Which I thought was interesting. It's like, oh, we have some unemployed hackers, do we? Oh, we'll give them something to do.

So as I noted at the top, to show potential affiliates that they mean business, REvil has deposited 99 bitcoins, approximately a million dollars, on the hacker forum. This particular hacker forum allows members to deposit bitcoins into a wallet hosted by the site. Members can see other members' deposits, and the deposited bitcoins can be used to privately buy and sell illicit services or data through the forum. Right? So it sort of functions as an exchange. And as of their posting, REvil now has 99 bitcoins deposited and on view in their hacker forum.

And what's interesting is that this deposit is also meant to demonstrate how much money ransomware operations are generating since they're publicly making a $1 million deposit as if it's no big deal.

Leo: Well, they're not giving it away, though. They're keeping it.

Steve: Well, except it shows they're not very concerned that the forum's administrators could steal it.

Leo: Oh, that's true, yeah.

Steve: Because they could. The hacker forums - and of course a hacker forum, as they say, there's no honor among thieves.

Leo: Welcome to the hacker forum. But you've seen, I think it's in Vegas, maybe it's in Reno, in Binion's Casino they have a million dollars under plexiglass. And it's the same exact thing, like see what you could win. It just gets your blood going. See what you could get. See what you could earn.

Steve: Yup. Yup. So unfortunately I think it is the case that ransomware is here to stay. I mean, it has set itself up as, I mean, remember back in the quaint days at the beginning of the podcast, Leo, when we were talking about email worms or email viruses. And it's like, oh, you know. You don't want one of those. And we were like, I wonder why anyone bothers to do that?

Leo: Yeah, it was just - it was like vandalism in those days.

Steve: It was a kind of a hobby to see maybe if you could get your name in the tech press or something. Now, thanks to bitcoin, there's money in them that malware.

Leo: That's a really good point, the idea that you can get anonymous - you would have to go to the 7-Eleven and buy a bunch of payment coupons.

Steve: Yeah, I mean, and remember it used to be like Western Union. The FBI would track down Western Union payments and nab people. And so it's bitcoin which has facilitated this dark industry because now it's possible to move money through a public blockchain, which we didn't have before.

Leo: Is it fair to blame the IT guys at these companies? Like it's not that easy to mitigate this stuff.

Steve: Okay. Which is, again, Leo, you are just hitting these segues perfectly today. Which brings us to UHS. Over this past weekend, starting last Saturday night, to better avoid detection, the huge international 400 healthcare facility strong Universal Health Services, which is currently 330th place on the Forbes 500 list of the largest U.S. publicly traded companies, was hit by a huge Ryuk ransomware attack. Early Sunday morning they shut down systems at healthcare facilities across the U.S. Based on reports from UHS employees, UHS hospitals in the U.S. including those in California, Florida, Texas, Arizona, and Washington, D.C., are left without access to computer and phone systems. And in some of the reporting - I meant to put it in here, but I forgot - there was a map of

these 400 facilities. I mean, they're all over the country. At the moment, the affected hospitals are redirecting ambulances and relocating patients in need of surgery to other nearby hospitals.

Leo: Wonder how computerized these facilities are.

Steve: Completely. They're completely reliant now on non-paper-based technology. One of the reports said: "When the attack happened, multiple antivirus programs were disabled and hard drives just lit up with activity. After one minute or so of this, the computers logged out and shut down. When you try to power back on, the computers just automatically shut back down. We have no access to anything computer-based, including old labs, EKGs, or radiology studies. We have no access to our PACS radiology system." Employees were told to shut down all systems to block the attackers from reaching further devices on the network.

Leo: It's a little late for that.

Steve: Uh-huh. And sadly, Leo, four deaths have now been ascribed to the incident due to doctors needing to wait for lab results to arrive via courier since all electronic channels are down.

Leo: This is a heinous crime. Murder.

Steve: Yeah. And reports are that during the cyberattack, files were being renamed to include the .ryk extension, the Ryuk ransomware hallmark.

Leo: Wow.

Steve: Based on some information that was shared with BleepingComputer by Vitali Kremez of Advanced Intel, the attack on UHS systems likely started through a phishing attack. According to Vitali, their Andariel, they have something they call the Andariel intelligence platform, detected both the Emotet and TrickBot trojans affecting UHS Inc. throughout 2020, and more recently this month. The Emotet trojan is spreading via phishing emails containing malicious attachments that install the malware on a victim's computer. After some time, Emotet will also install TrickBot, which ultimately opens a reverse shell to the Ryuk operators after harvesting sensitive information from compromised networks. Once the Ryuk actors manually get access to the network, they start reconnaissance. And after gaining admin credentials, they deploy ransomware payloads on network devices using PsExec or PowerShell Empire. In other words, the way this is being done is well understood. It's no mystery at all.

Leo: So they knew about it. Is there so much malware wandering around that you see a few Emotets, you go, yeah, there's just some more Emotets.

Steve: Yeah, exactly. And much like the APT, remember, the Advanced Persistent Threat that gave Sony Entertainment so much trouble all those years ago, I would argue, to answer your question, it's virtually impossible to secure anything that's as sprawling as a

400-facility network where everybody has PCs with email, and all of their workstations are tied in in order to be able to do their job. I don't know how you possibly make that secure.

UHS put out a statement to the public saying: "The IT Network across Universal Health Services facilities is currently offline, due to an IT security issue. We implement extensive IT security protocols and are working diligently with our IT security partners to restore IT operations as quickly as possible. In the meantime, our facilities are using their established back-up processes including offline documentation methods. Patient care continues to be delivered safely and effectively." Those four deaths notwithstanding. "No patient or employee data appears to have been accessed or copied or otherwise compromised."

So, you know, that was the standard publicly traded entity CYA notice. But there was that huge outage in the U.K. that brought down its health system. And here we have a massive, massive event.

Leo: They're probably attacking hospitals because of that precise issue, which is you've got a lot of different computers with a lot of different operating systems in a variety of different patch states.

Steve: In a hierarchy of training of employees. And remember, I mentioned it last week, there was an item that I didn't get into the show notes that talked about how studies have shown that "do not click that link" training fades after about six months. So like you put everyone through training, and then for a while people are being diligent. But then in comes that email from Aunt Martha, and it's like, oh, and you reach out and...

Leo: I guess no scanner can stop all threats. I mean, you try to stop malware from coming in via email at all; right?

Steve: Well, yeah. And the problem is the unknown threats. We'll be talking about Zerologon in a minute. That is now eight weeks old. And three instances of Zerologon malware were found uploaded to VirusTotal. Only a little over half of VirusTotal's scanners, like 73 of them. Today, because I reran the scans last night, they got like 43 out of 73 hits.

Leo: Yeah, yeah. I've said this all along. No antivirus is more than 50 or 60% accurate.

Steve: And that's the problem is that, because you're able to use encryption, because you're able to dynamically encrypt each instance of the malware so that there's no pattern, if you can't decrypt it when it goes through your scanner, as we've often said, proper encryption turns a file into maximum entropy noise. There is no pattern to be seen. So, I mean, it really just is - it's a scourge. It's nothing less than that.

Leo: Yeah. And there's also this whole thing about the delay, the advanced persistent threat thing, the delay between the time you get into the system, the system's compromised, and the time you trigger the ransomware, that time can be used to figure out what the backup strategies are, to figure out where the backups are stored and to put malware on those systems. I mean, if you have time - and

then of course, lately, exfiltrate information that you can then also use in the blackmail; say, well, maybe you've got backups, but we still have all your customer data that we're going to release. And we're seeing that happen a lot lately.

Steve: And it's easy to say, oh, well, why don't they have egress filtering?

Leo: Right, right.

Steve: They should be doing like bandwidth monitoring. Okay. How do you do that at 400 individual facilities?

Leo: You just don't. That's right. Yup.

Steve: I mean, some are going to be extended care facilities. Some are going to be, I mean, they're just a spectrum of different profiles. And something gets in, you know, and sits quietly on a machine in the corner and waits until the evening when no one's around, and begins to reach out and explore.

Leo: Explore. And they mentioned that radiology system, their PACS radiology system. Dollars to doughnuts that's running XP Embedded or something else ancient.

Steve: Uh-huh.

Leo: Right? Right. And you know it's not air-gapped or it wouldn't be down.

Steve: Right. Well, and Leo, it wouldn't be useful if it were air-gapped.

Leo: It wouldn't be useful, yeah.

Steve: Because those big CAT scans are huge files, high-resolutions, multi-sliced. They're big files. They have to be on the network in order to schlep those files around.

Leo: Right. Yikes.

Steve: No, I mean, it's...

Leo: It's terrifying. And it's unconscionable. I mean, the people who are doing this, obviously money means more to them than human lives. It's very sad.

Steve: Well, and they hate the U.S. I think it's fair to say.

Leo: Oh, you think it's also political? Yeah, maybe. I don't know.

Steve: Well, if not political, not in terms of like presidential race political probably. But, I mean, we know Russia is not a friend of the U.S.

Leo: Right, right.

Steve: And a lot of this ends up getting traced back to formally sponsored attacks. And it is weird, too, because when I read about like known Chinese state-sponsored attacks, I think, really? I mean, we know this? And we're doing nothing about it? The only thing I can think is that we're doing it against them just as much as they are against us.

Leo: Oh, I'm sure, yeah.

Steve: And so, well, in that case it's not like we're running around being Boy Scouts, and everybody else is throwing dirt at us.

Leo: I wonder if at some point there won't be a cyber Geneva Convention where we agree that it's just too dangerous to use these weapons on each other.

Steve: Well, like a form of the fact that nuclear arms have not been used since they were first used.

Leo: Exactly, exactly.

Steve: Because everyone said, okay, look.

Leo: It's mutually assured destruction.

Steve: Let's not all just nuke each other.

Leo: Yeah.

Steve: Yeah.

Leo: It's getting to be like that, where it's just escalating cyberwarfare. The long-term threats, even the short-term threats are not good.

Steve: Well, and everyone's society is becoming increasingly dependent upon it. I mean, it is - we're about to talk about what may not be incredibly useful technology, this flying spy cam. But most of it is really important. I mean, think about it. Everyone listening to

this podcast is able to think, I wonder what, blank, whatever. Fill it in. And yes, and go duck it. Or google it.

Leo: Oh, you can't say "go duck it." No one knows what - I was like, what is - is he having a stroke?

Steve: Exactly.

Leo: Go duck it. What? You're at DuckDuckGo. Okay. Go duck it. I like it. Let's do that from now on. Okay.

Steve: So any question we have can be answered.

Leo: Right. It's remarkable.

Steve: Just it's unbelievably powerful.

Leo: We live in amazing times, good and bad.

Steve: Yeah. Okay. So under our rhetorically named "What could possible go wrong?" section, we have Amazon Ring's announcement Thursday of their autonomous home security flying webcam.

Leo: Yu know, the thing that surprised me when they announced this, Lisa has forbidden cameras in the house forever.

Steve: Yes.

Leo: Rightly so. She doesn't know where that video is going. I said, yeah, look at this, you wouldn't want this. She's all, no, when are we going to get it? I'm like, oh. Because the camera's hidden in the base until it takes off.

Steve: True, true.

Leo: Right? But in order to use it you have to map your home. You actually have to take it around and say, okay, this is the garage. This is the kitchen. You have to teach it the layout of your house.

Steve: Yeah.

Leo: Like the first thing you do.

Steve: It's named the Always Home Cam. It will cost $250, and it's slated to start shipping next year. It is self-docking to allow it to recharge, and it's able to fly around its owner's home on pre-approved paths, as you mentioned. And it's supposed to allow homeowners to check to see if they left a window open, forgot to turn the stove off.

Leo: Yeah, I like that, yeah.

Steve: Or to check to make sure robbers aren't breaking in, presumably, because they're also billing it as a security feature. And perhaps not surprisingly, this announcement has been met with some mixed feelings.

Leo: Yes.

Steve: Rick Holland, the CISO and VP of Strategy at Digital Shadows, in an interview he told Threatpost: "For privacy advocates, the concept of an untethered IoT device surveilling the house is a little disturbing. Coupled with Ring's controversial privacy practices, the adoption of the drone could be low. However, those that have already embraced the concept of in-house security cameras are likely to be excited. The prospect of having a single drone monitor your house instead of multiple individual cameras would be alluring."

And this is exactly Lisa's position. It's like, hey, instead of having cameras steadily looking out into our various rooms, this thing gives you total coverage, potentially, as long as it's you who are flying it around and looking at things. Because consider the idea of it getting taken over and someone in Russia taking a stroll around your house using this thing. I mean, really, that's what's going to happen.

Ring, for its part, said that they built in privacy into the physical design, noting that when the drone is docked in its charging base, the camera is physically blocked. The camera's down on this - it's sort of like a - think of it as a big T-shape where the top of the T is a big square. So it's a square with a square post that drops down from the center of this T, and the camera is down at the bottom of that. So when this thing descends into its dock, the post is going into a square-shaped hole, thus preventing this thing from seeing anything. So it's clear that the camera cannot be seen.

But it's funny, too, because they said: "The device also has been designed to hum at a certain volume, so it's clear that the camera is in motion and recording." And I thought, wait. Hum? Are you kidding me? Has anyone here not ever heard a micro drone fly?

Leo: Drones are loud, yeah.

Steve: You can't hear yourself think. Generating the lift required using four tiny designer-approved props, because of course this thing has to look cute and right, that requires that they spin at thousands of revolutions per minute. At least I suppose we don't need to worry about the thing creeping up on anyone and surprising them. And then what occurred to me is what I want to know is whether no one who designs and tests these things has a dog or cat at home. Because this thing would drive the pets insane.

Leo: Yes, it would.

Steve: They would dive under the bed and never be seen again.

Leo: Terrifying.

Steve: There's this thing flying down the hallway.

Leo: You're right, I hadn't thought about that. Absolutely.

Steve: Oh, my god.

Leo: Wow.

Steve: Which brings us to another of this podcast's favorite aphorisms: "Not everything that can be done, should be done."

Leo: Yeah, yeah.

Steve: I had a neat link that I found posted to GRC's health newsgroup this morning that I wanted to share. And in fact we have time, Leo, if we could play just the first three minutes of this podcast, I mean, the first three minutes of this YouTube video into the podcast, I think that our listeners would enjoy it. The doctor speaking is a well-known U.K. M.D. I've seen some other of his YouTube videos. And I commend them to our listeners. But this is really good.

[BEGIN YOUTUBE CLIP]

DR. JOHN CAMPBELL: You are welcome to today's video. Thank you for coming back. Quite a few interesting new developments today. But I'm going to start off with an intriguing story about Vitamin D. Now, we mentioned a few days ago when we were looking at an interview with Dr. Anthony Fauci, the leading infectious diseases doctor in the states, he mentioned in the interview that he was personally taking Vitamin D and Vitamin C. So I was intrigued to know the dose because is it the same as the official guidelines was kind of my main question. And this is the story that's unfolded.

So this is from a Dr. Kari Hjelt, whose name I've probably pronounced wrong, doctor, sorry. Anyway, he actually wrote to Dr. Fauci. "You mentioned in a recent interview that you take Vitamin D supplement. For the greater audience, it would be most interesting to know the dose that you use." And as far as I can tell, this is completely genuine. It looks like a completely genuine National Institutes of Health email. And Dr. Fauci has written back and said 6,000 IU per day. Fascinating. And then Kari bounced that back on to me and has given me permission to share his name and details. So appreciate the initiative that you've taken there.

Now, so what this is saying is Dr. Fauci is taking, as far as we can gather from this, 6,000 IU a day, which is equal to 150 micrograms. So that's the dose he seems to be taking. Now, we have mentioned a few times that when it's not sunny, I am personally taking that.

Leo: Nice fountain pen.

DR. JOHN CAMPBELL: Fifty micrograms. So that's what I'm taking. Now, we have to stress, and I have to keep stressing this, I can't prescribe you a dose of Vitamin D for your requirement. So this is not me prescribing to you. It's not possible to prescribe on the Internet. But I am reporting what I take. Now, let's compare.

[END YOUTUBE CLIP]

Leo: I've been taking 5,000 per your recommendation.

Steve: And that's what I'm taking.

Leo: I think that's adequate. But it's really interesting to compare this to the RDA for Vitamin D, which is one, what 30th of that amount?

Steve: It's 800 at this point, 800 IU versus 5,000. And I think there's a broad misunderstanding about RDA and where it came from, you know, Recommended Daily Allowance, Recommended Dietary Allowance. The RDA was developed by the military back a long time ago to determine the content of military rations.

Leo: Oh.

Steve: They were trying to determine what they had to put in there to keep soldiers from dying from eating K-rations. And so people assume that that's, like, recommended, that is to say, like it's what you need to be healthy. It's not. It's what you need to keep from dying. It's the absolute minimum amount of things you need to take. And so, yes, you won't get rickets. You won't get scurvy. I mean, the RDA will keep you from becoming unhealthful. But there is a wide range between what keeps you from being sick and what you need for brimming health. And so anyway, the RDA is like, yes, everybody should get at least this. But often you should be getting way more than that.

Leo: Much more, yeah.

Steve: And that's really the case. And so yes, I just, you know, we keep seeing report after report after report that, especially now with COVID, it's so inexpensive, Leo, $15 for 360 little tiny drops of sunshine, little golden capsules, per person.

Leo: Well, also I note that the recommended allowance does go up after 71. And this is because, as you age, your ability to synthesize Vitamin D, which can only be done through the skin in sun, diminishes.

Steve: Well, and remember, in my famous Vitamin D podcast was that occurred after I had been naked sunbathing at noon and going to the lab daily to see whether I could develop any evidence of my own ability to synthesize Vitamin D. I did, I sort of blocked

off the patio so that I wouldn't be too - come to the attention of my neighbors. But I did that. And I think I did it for a month, and nothing happened. I got zero effect.

Leo: Yeah. It's a lot easier to take that pill. And there's other side effects to too much sun exposure. That's the other side of this is, of late, most of us are slathering on the sunscreen, wearing wide-brimmed hats. We're trying to stay out of the sun. And so even if you're under 71, I would guess that the required amount of Vitamin D orally has gone up quite a bit. Now, he said Vitamin C. By the way, I'm going to watch Dr. John Campbell. I'm going to watch him. He's great.

Steve: Oh, he is really...

Leo: I like his style. He's not doing the big YouTube, da da da, hit the button, smash the dough. He's just giving some straight facts, and he uses a very nice fountain pen.

Steve: He's a real scientist. Yeah, he posted something maybe a month ago, also on the subject. Maybe it was the genetics. I don't remember. But anyway, his stuff is really good. So I can commend him as someone worth browsing around.

Leo: Dr. John Campbell, and he's on YouTube. Now, he also mentioned something about Vitamin C. Does Fauci have anything to say about Vitamin C?

Steve: Don't know. That wasn't a question that was asked. But you're right, he did say both.

Leo: Said something, yeah. Because I'm taking, per your suggestion, a lot of Vitamin C. Not as much as you. But doing 3 grams a day.

Steve: I believe it's - I think it's another mistake that we're making by not having - people not having a lot more Vitamin C in their diet.

Leo: And the problem with Vitamin C is you can't store it, so it gets excreted immediately. So I'm titrating it. Like Lorrie, I'm drinking it in my water, dissolving 3 grams in my water and drinking it all day, so I get a little bit every few minutes.

Steve: Yeah. There's a strong case to be made is that we and fruit bats and guinea pigs are the only mammals that are not making our own Vitamin C. The household dog and cat, giraffes, elephants, zebras, tigers, I mean, all other mammals synthesize it because it is such a powerful antioxidant for dealing with the consequences of our metabolism. And in fact what I find fascinating is that our livers are trying to make it. It's made from glucose in a six-stage process. The glucose molecule is manipulated through six different enzymatic chains to come out as Vitamin C, where it then becomes useful. And our liver is doing all of the first five.

But the last enzyme, that's called L-gulonolactone oxidase, if anyone's curious, it's missing from the - we know where the defect is on the human genome where back in our

evolutionary past it got broken. And it must have been that our diet then was so heavy in Vitamin C that losing the ability to synthesize it endogenously didn't matter. We were getting it from diet. But now we're eating steak that doesn't have any Vitamin C in it. And so it's like, whoops. And the problem is anything that kills us quickly we find a cure for, or we focus on it a lot. It's the things that kill you really slowly that they just slip under the radar. So anyway, yeah.

Leo: Great stuff. I'm going to watch more of him. That's really...

Steve: Yeah, he's really great. I do commend him to everybody.

Leo: Dr. John Campbell.

Steve: This also just happened. Someone posted a link in the spinrite.dev group with a note that was found out in the wild, actually on Reddit, about the InitDisk utility that I created along the way to getting us to SpinRite. Remember that in order to use SpinRite you need to boot to DOS and then run SpinRite, which has always been and will continue for the foreseeable future to be a DOS utility. That requires that you be able to set up a thumb drive and make it bootable. There was no - I tried to go simple ways. I ended up having to go down to what I call "bare metal," where I'm manually writing sectors on the physical device in order to prepare it, to establish a FAT format, to establish the partitions. I did it all by hand.

But in the process what seemed to happen is that our initial testers were reporting that USB devices they had put in the bad bin, this thing was bringing back to life. So anyway, this posting was titled on Reddit, and I've got a link to anyone who has questions - actually it's under Linux questions under Reddit. It was titled "Stick is not working." The poster said: "I have formatted my stick with 'sudo cfdisk /dev/sdb.' And after I have formatted it, it has not worked anymore. Even after formatting it multiple times with cfdisk and gparted, it did not work. I've even tried it with Windows, but it also did not work.

"I've also tried multiple instructions from websites that also did not work. In gparted it says that the reasons might be, one, that the file system might be damaged; two, that the file system is unknown to gparted (which can't be true because I have formatted it with gparted multiple times)," he says in parens. "Three, there is no file system. Four, the /dev/sdb2 is missing." And then underneath it, it said in all caps: "PROBLEM SOLVED." Because down in the thread, the first person to respond, this was - looks like DoktorS_DE. So the first person to respond was roachh2, who said: "It's probably damaged." Then Patient-Hyena posted: "Try GRC's InitDisk. It is a Windows utility, but I have seen it work miracles." So then DoktorS_DE, the original poster, replies: "I will try that." And then he followed that up with: "It worked somehow. Thank you very much."

So again, just a reminder to our listeners just sort of, you know, it's free. And it is becoming sort of like the go-to solution for solving anything having to do with a USB device that won't boot. If this thing can't make it boot, then I would argue, you know, I can't bring it back to life, give it a file system that you can then do a directory of and so forth, it really is dead.

Leo: What file system does it write? What is it doing?

Steve: It puts a FAT32 on.

Leo: Okay.

Steve: And that's handy because of course that's universal. Everybody is able to understand FAT32.

Leo: Right, right. And I wonder how it overcame whatever was going on.

Steve: Well, it just is a big powerful plow, and it goes and plows its own furrows. It just doesn't care what's there. It's very careful about making - the way I had it operate is you start the program and then, while it's watching your USB, you plug in the device that you want it to kill.

Leo: Oh, interesting, oh.

Steve: Yeah. That way there's no, like, oh, wait, did I reformat the wrong one?

Leo: Which one, yes. Yeah, that's smart, yeah.

Steve: And so it takes a physical action from the user. And so it sort of walks you through the process step by step. And then it just says, okay, say bye-bye, and it just wipes out what was there and then reestablishes a new kingdom.

Leo: Windows only? DOS?

Steve: Windows only, yeah.

Leo: Windows only, okay.

Steve: And that's what this Linux guy, I mean, this was a Linux forum where the guy said, well, it is a Windows utility, but it may bring the thing back. And we know it did.

The new GRC web forums are now up and running. As we've been covering round after round of catastrophic WordPress add-on disasters, I've been growing more and more nervous about hosting my own public WordPress blog on my own servers.

Leo: Yes. As we know, that's risky.

Steve: Oh, lord. It's just been a disaster. It's difficult to be convinced of its safety without investing far more time in it than I want to. And I'm not a prolific blogger. This podcast audience gets pretty much everything I have to say, every week. And aside from being here with everyone, I work over in GRC's newsgroups. I only had a single blog

posting in years. And so, you know, the risk-reward ratio of hosting a WordPress presence on my own server was all wrong. If I were to do it again, I would go back to letting WordPress host my blog on their server and keep WordPress as far away from my servers as possible.

Anyway, I realized that I could create a blog forum, my own personal blog forum, on GRC's new forum system and have everything then integrated into one place. So that's what I've done. With great relief, after setting that up, I completely shut down and removed WordPress from my site and sight - S-I-T-E and S-I-G-H-T. Period. So as for the rest of the new web forums, we're still in the staging stage, awaiting the finalization and readiness of the ReadSpeed low-level driver-testing mass-storage benchmark, which I will now be returning to work on finalizing. I had to take a break in order to get the new forums all up and running, update my SPF DKIM and DMARC records so that I'm cryptographically signing outbound email. There's a lot of prep that goes into getting everything ready.

But I am ready to invite our podcast listeners who would like to register and claim their userID to do so. It's my intention that these forums will eventually grow to become GRC's and my own primary public presence for managing everything I do moving forward. So we're establishing this as a significant asset for GRC. There are still no links to the forum from GRC. It's sort of nominally non-public. And I even removed the forum's obvious registration UI, sort of to make it invitation-only.

The URL is simple. It's forums.grc.com/register. And so for our listeners, podcast listeners, https - ooh, I didn't have an "s" in my URL in the show notes, but it'll bounce you over to "s" if you don't - https://forums.grc.com/register. And everybody's invited to go grab themselves a userID. You don't have to register in order to get ReadSpeed or anything. But if you imagine you might eventually want an ID there, you can get one before it's gone.

Leo: Whose forum software are you using, just out of curiosity?

Steve: I'm using XenForo.

Leo: Oh, so it's the same as before.

Steve: Yes, it is. I really like XenForo, and I've gotten a lot of positive feedback from people. Before I shut down the blogs, I sent one final blog and let everybody know: Thanks for following me on the blog; I'm going to be over here from now on. Oh, and that's what I wanted to mention. Since we all hate email spam, I've configured the forum to never send any email notifications unless specifically requested. So I would imagine our listeners would probably want to click the "watch" button at the top of my blog forum in order to receive a note when I post news there, as I certainly will be. Although listening to this podcast you pretty much know what I'm up to anyway.

Leo: Just got my name.

Steve: Cool, yay.

Leo: Can't let LeoLaporte because I got it. So if I had an account, as I did on the old forum, it doesn't propagate over to the new one.

Steve: Correct.

Leo: Got it.

Steve: Does not cross over.

Leo: Okay.

Steve: And I did get a bunch of people because first I opened it to the newsgroup denizens who have helped bring me to the point that we're getting ready to launch this software, and we'll be doing the same thing for SpinRite. Many of them were SQRL users. And of course if you did present your SQRL identity, you're just instantly registered. I mean, many of them said, my god, this is the smoothest registration process I've ever seen.

Leo: Yeah, it's nice. Really well done, yeah.

Steve: So anyway, all that stuff is working, yeah. So we're set and ready. I will now go back to work on getting the software finalized, and I will be announcing it on my blog there and on the podcast when it happens.

Leo: Nice.

Steve: Okay. So as we know, last week one of the many topics that we covered was the so-called Zerologon vulnerability. This week it has grown to become this week's main subject, for a good reason. Last week I introduced us to Secura's earlier discovery and their responsible disclosure which resulted in a quiet patch as part of August's Patch Tuesday. Secura waited six weeks, feeling that that was appropriate, giving them two Patch Tuesdays' lead. And then they released their full description of the vulnerability. Although they had created a working proof-of-concept demo, they deliberately chose to withhold it even then, six weeks downstream, to allow more time to pass for the August updates to be more widely deployed.

But while withholding the proof of concept was a nice gesture, the lesson to the industry now going forward is going to have to be that all anyone needs to turn a nebulous vulnerability description into an exploit is a sufficiently clear description such as Secura provided. In other words, them not publishing a proof of concept didn't matter because within hours, and I mentioned this last week, of Secura's disclosure, multiple working Zerologon proof of concepts had been created, worked out, tested, and verified, and were there on GitHub.

One week ago, when I first talked about this, there were three publicly posted exploits. Today there are more than the first page of results returned by Google. I did a search. Yes, I didn't duck it, I googled it. We've got on GitHub from RiskSense Zerologon exploitation script. We've got from Zero Networks Zerologon test for SMB and RPC. Oh,

and we'll be talking about that because you mentioned that last week, Leo, about Samba. We've got the full coverage of that now. This says "demonstrates that CVE-2020-1472 can be done via RPC/SMB and not only over RPC/TCP." So they have something.

Then there's Invoke Zerologon. There's the Zerologon testing script. There's a 1472 proof of concept. There's something called Zer0Dump, saying it's a proof-of-concept exploit tool for abusing the vulnerabilities associated with CVE-2020-1472 Zerologon in order to initiate a full system takeover of an unpatched Windows domain controller. We've got the Zerologon exploit, tests whether a domain controller is vulnerable to the Zerologon attack. If vulnerable, it will reset the domain controller's account password to an empty string. Then we've got one just called CVE-2020-1472 from VoidSec, tests whether a domain controller is vulnerable, blah blah blah. Same thing.

Silverfort's scanner for vulnerable domain controllers with Zerologon. There's a scanner there. We've also got from Rsmudge Zerologon BOF. We have Add Exploit, adds an exploit module - ah, this is Metasploit - to the Metasploit framework. He says this is a pure Ruby implementation leveraging the changes proposed. And then there's a Ruby SMB deal.

So what started off as a few is now an explosion because this is not difficult to do. And as we know, when they are easy to do, so it doesn't require climbing a high mountain, and there is dramatically large potential payoff and a huge potential target base, all of those conditions are met, the result is attacks. Of course then we know that when I was talking about this first last Tuesday, it was just the prior midnight was the deadline for the DHS/CISA's emergency directive requiring all federal agencies, period, stop what you're doing, we don't care, no excuses, you have to apply this patch before midnight Monday night.

And, finally, the day after that podcast, so that was last Wednesday, Microsoft Security Intelligence account tweeted: "Microsoft is actively tracking threat actor activity using exploits for CVE-2020-1472 Netlogon EoP" - they're calling it an elevation of privilege, which I think is cute - "vulnerability, dubbed Zerologon. We have observed attacks where public exploits have been incorporated into attacker playbooks."

So of course from GitHub there they are, and they are immediately grabbed up by the malware guys and turned into active attacks. And I note that their use of EoP, Elevation of Privilege, you know, I guess technically the Zerologon vulnerability is an elevation of privilege. After you subvert the connection's encryption by using null initialization vectors, bypass the domain controller's authentication, null its password, and then logon as the domain admin, yeah, I'd say that you would have definitely elevated your privileges in this instance. That sequence, as Secura put it, was one, spoof the client credential. Two, disable RPC signing and sealing. Three, spoof a call. Four, change a domain's AD password. Five, change the domain admin password. And all that happens in a split second.

As I mentioned before, there were three files. Microsoft Security also tweeted three hashes for three files which are known to be used in this exploit. BleepingComputer tracked those three down by their hashes to the samples that had been uploaded to VirusTotal. All three were named SharpZeroLogon.exe. I've got the three links in the show notes. And they show, respectively, 45 out of 71 detections, meaning that as of last night - because as I mentioned I refreshed them when these detection rates still seemed surprisingly low to me.

I mean, like - and there were some worrisome major AV providers that are still not seeing this - 45 out of 71 detections for the first, 42 out of 69 for the second, 43 out of 71 for the third. So with those detection rates, certainly this raises those files well above any sort of false positive. But they still seem really low, given the severity of the

instance. I mean, like a lot of people are seeing it. But if you go click on those links, and I didn't bother to go through it all, but there are some major AV packages that are not finding it. And I don't think I'd want to be relying on any of those AV tools that today, eight weeks later, were still blissfully unaware of the Zerologon threat at this stage.

BleepingComputer wrote: "In one of the samples examined by BleepingComputer, and like other public exploits, the NTLM" - NT LAN Manager, which is what we're still using today - "the NTLM hash of the domain controller will be changed to" - and then they provide the hex, it starts out as 31 dog 6 charlie fox easy 0 dog 16 able easy 93 and so forth - "is the hash for an empty password." So when you look at the domain controller's hash, it's like, whoopsie, that's a null password.

Okay. So now for the Samba connection which was just coming to light last week. It turns out that Samba, the open source implementation of Server Message Blocks, or SMB - thus the name Samba. You add an "a" between the "S" and the "M" and an "a" on the end, and you've got Samba from SMB. That's of course the SMB protocol available for Linux and Unix systems, which maps directory shares between Windows and the Unix and Linux systems and incorporates the Netlogon protocol; and, as a consequence of also offering Netlogon protocol, suffers from the vulnerability.

Red Hat had a good write-up about this from their perspective. They said: "Red Hat is responding to a vulnerability" - and there we have again 2020-1472 - "in the Microsoft Netlogon service. Netlogon service is an authentication mechanism used in the Windows Client Authentication Architecture which verifies logon requests; and it registers, authenticates, and locates domain controllers. The Netlogon service, as part of the domain controller functionality, implements Microsoft Netlogon Remote Protocol. The implementation of Netlogon protocol contains a flaw that allows an authentication bypass. This was reported and mitigated by Microsoft as CVE-2020-1472. Since the flaw is a protocol-level flaw, and Samba implements the protocol, Samba is also vulnerable.

"The Microsoft Windows Netlogon Remote Protocol (MS-NRPC) reuses a known, static, zero-value initialization vector in AES-CFB8 mode." Which of course we talked about last week. "This allows an unauthenticated attacker to impersonate a domain-joined computer, including a domain controller, and potentially obtain domain admin privileges." They said: "In Windows environments, only the domain controller runs the Netlogon service accessible by clients. This applies to Samba when it is used as a domain controller. Samba Domain Controller role is implemented in both Active Directory mode and also the classic NT4-style mode. The Red Hat Enterprise Linux (RHEL) version of the Samba package only provides classic NT4-style domain controllers. An unauthenticated attacker with network access to a domain controller can impersonate any domain-joined computer, including a domain controller. The attack can result in a denial of service and potentially allow an attacker to gain domain admin privileges."

Anyway, I'll skip some of the other stuff. They get down to: "The Samba suite supports secure channel establishment between domain members and domain controllers. However, default behavior for secure Schannel prior to Samba 4.8 was to automatically negotiate secure channel only if a client supports it." And we know what that means. That's the classic security downgrade attack where the client pretends not to support the security, in which case Samba 4.8 says okay. Anyway, they said: "Since v4.8, the default behavior of Samba has been to insist on a secure channel for all clients, which is a sufficient fix against the known exploits of this attack. This default is equivalent to having 'secure schannel = yes' in SMB config."

Anyway, I'll skip the rest of this. Basically what this means is that, since v4.8 of Samba, which was released in March of 2018, so 2.5 years ago, if you have 4.8 you're fine because it insists on a secure channel to be established. Any earlier versions of Samba that haven't had their smb.config file changed to say "server schannel = yes," or if it still

says either "no" or "auto," those are vulnerable by default. So that's what somebody wants to do. If you've got actually any version of Samba, forget about version numbers, you'll know if you do. You should have "server schannel = yes," regardless of version, and you're fine because, if that requirement is enforced as it then would be, you have no vulnerability problem.

The reason there could be a side effect is there could be some devices which we talked about last week that just cannot do Schannel negotiation. So this is why Microsoft's fix at the moment is kind of a half step. They will be fully enforcing it the beginning of next February. For now, they're partially enforcing and logging. The logging portion then allows IT to identify those devices that are opting to use non-Schannel, and then they can go see about fixing them and updating them.

But that's how Samba got involved is that Samba did a straight-across, okay, this is the protocol that Microsoft originally created for their file and printer sharing for NT LAN Man. We're going to implement an open source version, and we're going to call it Samba. And unfortunately they inherited the problem.

There is a cool fix. Mitja Kolsec is CEO and co-founder of our friends 0patch, the people who rapidly step up and patch things that Microsoft hasn't patched or in some cases won't patch. They have another of their terrific micropatches. Sometimes they're a couple bytes long, which is just like you've got to shake your head. This one was a whopping 29 bytes because they decided they were going to go back and fix Windows Server 2008 R2.

Anyway, I'm getting ahead of myself. Their announcement blog post states, that is, the 0patch guys, it's numeric 0patch.com: "The Zerologon vulnerability allows an attacker with network access to a Windows domain controller to quickly and reliably take complete control over the Windows domain. As such, it is a perfect vulnerability for any attacker, and a nightmare for developers." In fact, I've heard now that this perfect vulnerability is sort of making its way through the industry because this is.

So in an interview with Threatpost, Mitja explained, he said: "Our micropatch was made for Windows Server 2008 R2, which reached end of support this January and stopped receiving Windows updates. Many organizations are still using this server, and the only way for it to get extended security updates from Microsoft was to move it into the Azure cloud, which," he writes, "is an unacceptable option for most organizations." He also added that 0patch is also working on porting their micropatch to various still-supported Windows servers for customers who, for various reasons, can't apply Microsoft's patch.

Okay. His blog provides a truly wonderful look into the micropatch development process, and it is this week's GRC shortcut URL, so https://grc.sc/786. Or you can also duck it by searching for "0patch Zerologon." And yes, I did just say "duck it."

Leo: There you go. There you go again.

Steve: And you can find it. Anyway, in his blog posting, this blog posting that I'm referring to, grc.sc/786, or just again search "0patch Zerologon," he explains in detail what's going on, what the patch developers faced, and what they created. And they include the full source code right there in the blog posting for the micropatch, which eliminates this vulnerability. They essentially duplicated what Microsoft did, that is, they looked at the change that Microsoft made, and they said, okay, we can do that.

The problem is, since the target platform is 2008 R2, meaning that it's 12 years old, the specific function that was patched by Microsoft in the later servers is not present. So they

had to implement the same patch functionality somewhere else in the authentication flow, and they did. But you're showing on the screen right now, Leo, just a beautiful documented reverse engineering of the problem that Microsoft had, and what they did. So the perfect vulnerability has exploded onto the scene. It's simple to do. It's incredibly powerful and destructive, if it can be used. It's every ransomware villain's dream come true, and GitHub is filled with sample source code implementations.

Leo: Boy, these patches are really tiny. I mean, this is, what, 30 assembly lines.

Steve: Yes. It's 29 bytes of code.

Leo: Instructions, yeah. It's nothing.

Steve: It's just beautiful.

Leo: Yeah. Wow. It's nice to see the code. I mean, I'm always nervous when you recommend a patch from a third party and all that. But if you can see the code, I guess that's okay.

Steve: Yeah, these guys have also - at some point reputations really start to count.

Leo: Yeah. I mean, they've been doing this for a long time, yeah.

Steve: Yup. And when they're providing a service that nobody else offers, I mean, you've got to say here's a third party patching an egregious problem in a server that Microsoft has chosen for corporate policy and commercial purposes to abandon.

Leo: Right. Right.

Steve: You know, again, it's difficult for that to be okay. But that's the world we live in, and we now live in a world where a lot of source code for this nightmare is out and is being quickly weaponized.

Leo: Got to do something, yeah. They've got to do something.

Steve: And Leo, it is every ransomware villain's dream, this thing, because it's corporations that are going to let, you know, bad guys are going to get in this way, install backdoors, and just have at it.

Leo: Probably already have.

Steve: Yeah.

Leo: In about two or three months we'll start to see it. That's Steve Gibson right there. This guy's doing God's work, as you can see. Microsoft's work, too, apparently. GRC.com is where to go to catch the latest version of the show. He has 16Kb and 64Kb audio versions of it. He has transcripts. He's also got SpinRite, the world's best hard drive recovery and maintenance utility available now in 6.0; 6.1 is almost here. You will be in on the beginnings, the beta tests, and the final release if you buy SpinRite 6 right now. Don't forget the forums, forums.grc.com/register. Early entry for people listening to this show. Get your name now. And they're looking nice, looking very clean, very sweet.

We also have audio and video of the show at our website, TWiT.tv/sn. Or you could simply watch it on YouTube or listen to it on your Amazon Echo. Just ask for the Security Now! podcast. Or better yet, subscribe in your favorite podcast client, and that way you'll just get it automatically the minute it's available.

We'd better get you out of here. I think you've got a date in about 40 minutes.

Steve: I'm watching the clock, Leo.

Leo: Me, too, Steve. That's Steve Gibson. Thank you, Steve, we'll see you next week on Security Now!.

Steve: Thanks, buddy. Well, you're going to see me in two days.

Leo: Thursday. Can't wait. 3:00 p.m. Pacific, 6:00 p.m. Eastern, our very special panel of Last Pass, no, ITProTV. Yeah. Thanks, Steve. See ya.