SECURITY NOW!

Transcript of Episode #899

Freebie Bots & Evil Cameras

Description: What happens when you run a Caller ID spoofing service? Or when you mislist and underprice online goods? Or click on a phishing link for a cryptocurrency exchange? Or consider working for a underworld hacking group? Use a web server from the Dark Ages in your IoT device? Or rattle your sabers while attempting to sell closed networking systems to your enemies? Or decide whether or not to continue to suspend your Twitter ad buys? Or log into Carnival Cruises with a Passkey? Or use hardware to sign your code? This week's podcast answers all of those questions and more.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-899.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-899-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with a potpourri of security stories: the end of a famous Caller ID spoofing service taken over by the feds now; a funny little scam involving misplaced decimal points; a web server from the Dark Ages that's unfortunately still being widely used; and when Passkey is not really Passkey. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 899, recorded Tuesday, November 29th, 2022: Freebie Bots and Evil Cameras.

It's time for Security Now!, the show where we cover your security and privacy online with the hero of the hour, Mr. Steve Gibson. Hello, Steve.

Steve Gibson: This is the podcast which just has a boring start every week because everything works. We're not spending half an hour trying to get, like, stuff onscreen or the lighting right.

Leo: Oh. You just don't see that part. In the old days we did that all the time, didn't we.

Steve: Yeah, I mean, well, and I mean, like with different hosts and juggling things and all that. But it is easier to do a one-on-one podcast than the thing you're doing with everybody else.

Leo: Oh, I see. You saw the rocky start of our previous program.

Steve: Yeah, and it's generally like half an hour before things...

Leo: Yeah, I know, it's a little weird.

Steve: ...get going.

Leo: It's a little weird, I know, I know.

Steve: Speaking of which.

Leo: Yes.

Steve: This is Episode 899 for...

Leo: Oh, dear. Oh, dear.

Steve: Oh, I know. And this is the birthday episode, for those who don't know. Leo is celebrating number 66.

Leo: Look at that. Route 66.

Steve: Route 66 on his - for those who don't have video, he just held up an old sign that won't mean anything to anyone much younger than us.

Leo: "Kookie, Kookie, lend me your comb." That's all I have to say, yes.

Steve: So this was one of those weeks where nothing really stood out, but a lot of interesting things happened. So I grabbed two of the items we're talking about as the title, basically taking from the typical naming of your other podcasts, Leo, where you think, okay, what did we talk about? Let's come up with something about that. So this is Freebie Bots and Evil Cameras for 899. And during this podcast we're going to answer a few questions. What happens when you run a Caller ID spoofing service? Or when you mislist and underprice online goods? Or click on a phishing link for a cryptocurrency exchange? Or consider working for a underworld hacking group? Oh, no, this is a great podcast. Or use a web server from the Dark Ages in your IoT device?

Leo: This is not all one story. These are multiple stories.

Steve: Oh, yes. Oh, yes, yes, yes. Good point. Otherwise that would be one hell of a story.

Leo: Really confusing, yes.

Steve: Yeah. Or rattle your sabers while attempting to sell closed network systems to your enemies? Or decide whether or not to continue to suspend your Twitter ad buys? Or log into Carnival Cruises with a Passkey? Or use hardware to sign your code? This week's podcast answers all of those questions and more.

Leo: Now, that's a tease. You are absolutely now, finally, after 899 episodes, conforming to the TWiT way.

Steve: It only took 17 years, my friend.

Leo: We were heading your way, and you headed our way. So we've met in the middle.

Steve: And then we're going to defrag a zebra. Just, you know...

Leo: Oh, wait'll you see this. That is our Picture of the Week.

Steve: That's pretty good.

Leo: And a good one it is.

Steve: That's right.

Leo: I feel like this is almost a Dad joke, this Picture of the Week.

Steve: Well, and I think we've used it before [Episode 637]. I mean, it looks familiar to me. But, if so, it's kind of fun anyway. So for those who are not seeing our video stream, the caption on this...

Leo: [Laughing]

Steve: I know, it's really good. It says "I defragged my zebra." And what we have is what looks a horse with the front half black and the rear half white.

Leo: It's defragged.

Steve: It's defragged, yeah. All of the in-use clusters got pushed to one end, and the free space is on the other. And anyway, it's just very clever.

Leo: I love it. It got us talking about defragging. You don't really do that anymore; right? Most modern operating systems that's handled.

Steve: Correct. Well, actually, the argument that Microsoft has always made, although it's not really as true, is that there is no - they were always saying there was no need to defrag NTFS file systems. It was clear that over time FAT32 file systems became fragmented. And what we were saying before we got on the air was I was posing rhetorically the question, how many user centuries of time were lost...

Leo: Watching.

Steve: ...with us just staring at the defrag screen, while the little squares jumped around. It was just wonderful. And, I mean, it served no constructive purpose whatsoever. But, you know...

Leo: We loved it. It was fun.

Steve: It just really - maybe it was a way for a geek to have a timeout. It's like, okay.

Leo: It was a zen. It was meditative, exactly. Yeah, yeah.

Steve: Yeah. Yeah.

Leo: But not necessary anymore. That's right; right?

Steve: Well, okay. So Windows says that it defrags, like, automatically...

Leo: In the background, yeah.

Steve: ...in the background. Which may be the case. The one place it can be useful is for data recovery. If your files have been defragmented, and you lose somehow, some catastrophe, the entire metastructure of your file system, and you really desperately have to have some file back, but basically if you've lost all of the metadata, there's no directory hierarchy, no directories, anything, somewhere out on your drive is a blob of space that a file occupies. And if it's contiguous, if it is defragmented, you can find it.

Leo: Right.

Steve: I mean, it's there in whole. But if it itself is scattered all over the place, and it was dependent upon the file system's pointer structure in order to reconstruct that file on the fly, you're really not going to be in such great shape. So, you know. But in the old days the reason we of course defragmented was because seek times were so long that, if pieces of a file were scattered physically around a drive, the drive's head would have to go jumping back and forth around, in and out on different tracks, grabbing little pieces of the file in order to get the whole thing. If the file was defragmented, the head would just go to the beginning and just maybe tick over sequentially a few tracks, depending upon how large the file was, but so it was less wear and tear on the drive because it wasn't having to jump all over the place just to get one file read, and it was a lot faster because you weren't embedding all these seeks in the middle of a file read.

Leo: Of course there's zero seek time on SSDs, so you...

Steve: Right. And so that's what changed is when we went to solid-state, suddenly all of that head-seeking disappeared, and it made no difference in terms of performance.

Leo: Although Microsoft quite cleverly, I think, instead of defragging SSDs, if you issue the defrag command, because you still have a defrag, I believe you still have defrag...

Steve: Yeah.

Leo: Microsoft says, yeah, well, we'll just trim the SSD. It's a way to invoke trim. And so Allyn Malventano always said, you should still be defragging because you're now trimming your SSDs. Although I think modern SSDs do trim, as well, in the background. It's kind of...

Steve: Yeah.

Leo: That's kind of necessary to keep the speed up.

Steve: Yeah, well, it's actually an OS-level thing because the SSD has no knowledge...

Leo: Oh, it doesn't know.

Steve: ...of the date - right.

Leo: But it could be in the controller. I thought it maybe was in the controller.

Steve: No, it's got to be in the OS.

Leo: No? Oh, okay.

Steve: And so the idea is that the drive itself has no knowledge of the file system. It's file system agnostic. But all of the operating systems now, Linux does it, and Windows does it. In fact, it came up relative to SpinRite recently because, if you were to do a write-level, which is level 3 or 4 in SpinRite 6.1, that leads the drive to believe, the SSD to believe, that its entire space is now in use because when you write to something, basically it flags that area as in use.

So what you can then do is, under Windows, there is a way to say, please trim this drive. And under Linux it sort of does it more easily, but you're also able to force it. And so that is one sort of a power user tip that we'll be getting to at some point with SpinRite is, once you do something on an SSD that writes to the whole thing, you then need to put it back into the operating system to let the OS say, okay, calm down here. These are the areas that are actually in active use, and all the rest of this, no. That's just completely free. And the point is...

Leo: It's hard drive garbage collection. We've been talking about memory garbage collection.

Steve: Yes, very much.

Leo: It's smart drive garbage collection.

Steve: Yes, yes.

Leo: Yeah, yeah.

Steve: Okay. So I asked a question at the beginning of the show: What happens if you run a commercial Caller ID spoofing site? Well, your site turns into the top of this podcast, it's on the second page here. Yes, and anybody who's interested can go there now, I went there yesterday, I presume it hasn't changed. iSpoof.cc is the domain name. I-S-P-O-O-F dot C-C. And what you get is a big page that says, "This website has been seized," and the various emblems of global law enforcement. And it says: "This domain has been seized by the Federal Bureau of Investigation and the United States Secret Service in accordance with," blah blah blah blah blah. Anyway, and then we've got Europol and London City Police and Cyber Police and all kinds, you know, everybody's involved.

Leo: Wow, yeah.

Steve: So, okay. Get a load of this interesting bit of happening: Europol and law enforcement agencies from several countries, including the FBI, have seized the servers and web sites of iSpoof, which was a service that allowed users to make calls and send SMS messages using spoofed identities. And Leo, if you were curious - oh, actually I have a link on the page below to the web archive Wayback Machine of iSpoof from before it was seized. And it's quite interesting.

Anyway, so the service launched in December of 2020 and advertised itself as a way for users to protect their phone numbers and identities online. But Europol said that iSpoof was widely abused - yeah, no kidding - for fraud because it allowed cybercrime gangs to pose as banks and other financial organizations.

An investigation into iSpoof began in 2021 after Dutch Police identified the service during one of its fraud investigations. The Dutch Police said they linked the service to a web host in Almere, where they deployed a wiretap that allowed them to map the site's reach and learn the identities of its registered users and administrators. Officials said iSpoof had more than, get this, 59,000 registered users before it was taken down just earlier this month.

U.K. Metropolitan Police said that 142 suspects were detained throughout the month of November. So they did a big sting operation globally, with more than 100 individuals detained in the U.K. alone, including iSpoof's administrators. Europol said iSpoof was being used to place more than one million spoofed calls each month, that administrators made more than 3.7 million euros, and that the service has been linked to fraud and losses of more than 115 million euros worldwide.

The U.K. police said they plan to notify all U.K. users who received spoofed calls made through iSpoof, which is nice of them. So anyway, as I said, I was curious to see what the site looked like before the global takedown which displayed that site seizure page, above. So I turned to the Internet Archive Project's Wayback Machine. And what I found was just sort of, you know, headshaking.

The top of the site's very modern-looking home page, which sort of has a floating iPhone there on the right, proclaims: "Protect Your Privacy with Custom Caller ID." And it says: "You can show any phone number you wish on call display, essentially faking your caller ID." And then down in their features they said: "Get the ability to change what someone sees on their caller ID display when they receive a phone call from you. They'll never know it was you. You can pick any number you want before you call. Your opposite will be thinking you're someone else. It's easy and works on every phone worldwide!"

So, yeah, you can imagine that all kinds of bad people with ill intent would be abusing this thing. I mean, like, you know, ex-boyfriends or stalkers or spouses or whomever, you know, whose calls you are not accepting would just figure out whose call you were accepting and then spoof it in order to get you to answer the phone. I mean, it's awful. Anyway, we've talked a lot about how insecure all of this is. The, what is it, SS7, the current Signaling System 7, is still allowing this to go on. I finally gave up and disconnected, actually I had three, I had a fax line and two landlines, because all I was ever getting was just junk calls. They were just, you know, it was awful.

So for me, the most disturbing thing about this story is that the site was up and running for nearly two years before it was brought down. You know, that was a ton of damage to be done. And, you know, you can imagine how the word-of-mouth of this was no doubt spread among the world's shadier types as this thing was allowed to continue.

So for what it's worth, I hope there are not alternative sites that are already up and going. I would be surprised, frankly, if there weren't. I should have done a google and looked around. It didn't occur to me until just now. But still, just sad that it took that long to get this down. And we're hearing about the encryption and the tightening of the intercarrier communications. It's one thing for a carrier to be secure within itself. But it's the gap between carriers where we need security. And, you know, they're just not in a hurry. It's like, why, you know, we have to make them do this. And so far that hasn't happened.

Okay. What is a Freebie Bot, you ask? A new class of bot has been identified. And this one does something that would be difficult to predict; but, once you hear what it does, you think, huh, is that illegal? Last Tuesday, the anti-bot research and security provider Kasada, whom we've spoken of before, shared the results of their latest threat intelligence, which detailed the growing prevalence of so-called "Freebie Bots." Freebie Bots automatically scan and scrape retail websites searching for and purchasing mispriced goods and services, purchasing these discoveries at scale before the error is found and fixed.

Get a load of this. Kasada research has found more than 250 retail companies recently being targeted by Freebie Bots, with over seven million messages being sent monthly - monthly - within freebie communities. Okay. Now just to be clear...

Leo: But this isn't illegal. Right? No. This is capitalism, baby. You screwed up.

Steve: That's right. So just to be clear, these are not Furry communities, these are Freebie communities. Nor are they Furby communities, but that's something else. Members within one popular freebie community used Freebie Bots to purchase nearly 100,000 products in a single month with a combined retail value of $3.4 million. But Kasada's research revealed that, due to significant underpricing, the total purchase cost of the goods for the Freebie Bot users was $882. This allowed some individuals to realize a monthly profit of over $100,000. Top items purchased using Freebie Bots during this period of time included off-brand sleeveless halter neck mini dresses; get this, Apple MacBook Air laptops; and deep cleansing facial masks.

Leo: It's an interesting Venn diagram. I don't know.

Steve: That's right. What's your overlapping customer matrix? Many pricing errors were the result of a decimal point misplacement, granting discounts as large as 99%. Using the speed and scale of a bot attack to rapidly purchase as much stock of these erroneously priced goods as possible, actors then turn around and resell the goods at the price they should have been, reaping a large profit.

So you can see how this could happen; right? Someone keying in a new item's retail listing gets into the habit of entering a decimal point before the last two digits of the price. But then they encounter a price formatted as a whole integer number of dollars without any cents, and without thinking they place a decimal point before the last two digits, thus inadvertently reducing the listing's price by a factor of 100. It turns out that, at scale, across the entire Internet, these mistakes happen enough to have spawned the creation of a new class of bot, automated retail mistake-finding bots, which will instantly purchase as much of something that's been mispriced as they're able to. So human ingenuity knows no bounds. I suppose that while this might not be technically illegal, it certainly is unethical and dishonorable.

Leo: Is it? Is it?

Steve: Well, you know...

Leo: No. I'm buying it at the listed price.

Steve: You know when the MacBook Air is offered for 50 bucks...

Leo: Not my problem.

Steve: There's something wrong.

Leo: Not my problem. That's a good deal. I'll take it.

Steve: How many can I have?

Leo: I guess it depends. If this is happening to, you know, your local Goodwill Store, that's terrible. And that's probably more likely where it is. Apple probably never makes a mistake like this because they have good software. But still.

Steve: You're right. It's probably taking advantage of people who can't afford...

Leo: Of small retailers, yeah, yeah, yeah. I mean, Apple's never going to misprice its Apple gear on its site.

Steve: I have seen oddly priced things on Amazon. You probably have, too, Leo.

Leo: Sure, all the time, yeah.

Steve: Where it's just like, what? That can't be right. You know? And I just, you know, I mean, it's for a left-handed screwdriver, so I don't need one. But still.

Leo: I mean, I'm the kind of guy, and I know you are, too, that probably would go, oh, that's a mistake. I'm not going to take advantage of that. So maybe it is unethical. I wouldn't do that. But still, depends I guess on the size of the company. The problem is...

Steve: As I said, once you hear the idea, no one is surprised.

Leo: Oh, it happens all the time, yeah, yeah.

Steve: Well, no, I mean, that a bot has been created to go scan for these mistakes in real-time.

Leo: Oh, yeah. Absolutely.

Steve: And buy up the inventory. Wow. Okay. We have the anatomy of a real-time cryptocurrency heist. The group PIXM Security, whose business is to protect end users from credential fraud, recently blogged about the details of an attack group they've been monitoring. The lengths this group will and does go to, to circumvent, like, one of the newer protections, the deliberate "authorized device" protections we're beginning to see more and more, where like if use a new device, you log in with somewhere you haven't logged in before, there's like, whoa, we haven't seen this device before, so we're going to jump you through some extra hoops.

So, okay. What's interesting here, and I think you're going to find this really interesting, Leo, is their report in detail of what's behind a true real-life phishing exploit. So, okay. And just to give you a hint, scammers will use in-browser chat to initiate a remote desktop session on a victim's device, approve their own device as valid to access the user's account, then drain the cryptocurrency from their wallet or wallets.

So, okay. Here are the details behind this. When PIXM's Threat Research team first started tracking the group, they were only targeting Coinbase, right, like the premier exchange. Then over the past month the group has increased their coverage, that is, the bad guys have increased their coverage to add "support," if you call it that, for MetaMask, Crypto.com, and KuCoin (K-U-C-O-I-N), in addition to Coinbase. So now four. The spoofed domains are the typical slightly misspelled, in this case subdomains of azurewebsites.net. So that's the hub of where they are. And so it'll be like, you know, conbase.azurewebsites.net or something like that.

The group employs working effective second-factor relay interception when a user is spoofed into going to a lookalike site. Regardless of the credentials the user enters, whether they're legitimate or not, since the spoofing site cannot determine that initially, the user will be moved to a two-step verification page after clicking "Login," where depending upon the platform in question they'll get what they're expecting, which is either prompted for a second-factor code or their phone number is prompted and used then to receive a two-factor code.

The criminal group will first attempt to relay the credentials they've been given and second-factor codes to the legitimate login portal which is associated with the platform they're spoofing. Once the user clicks "verify," they will be presented with a message, no matter what happens, telling them unauthorized activity has occurred on their account. Well, it turns out it's true, actually, but this is the bad guys trying to reel them in further. As with the original Coinbase attack which this group started with, this will initiate a chat window to keep the user on the phishing page in the event the two-factor code should fail, which of course the bad guys don't know yet because they'll get prompted for that after they attempt to log in, and should the threat actor need to start a remote desktop session with the victim to continue with this attack.

PIXM wrote that, in their experience, regardless of whether the victim enters legitimate credentials or not, the group will "chat" with the victim to keep them in contact should they need to resend the code or proceed to the second phase of the attack. The criminal gang's willingness to do this significantly increases, I'm sad to say, end-user engagement, you know, and their belief that, like, they're talking to the real guys, right, because there's someone there. For the majority of the attacks which this group carries out, they engage in direct interaction with the user. Their spoofed login and verification portals will, by default, return a login error, as I mentioned, regardless of the actual standing of the user's account on the actual exchange and the wallet.

Of course this process is intended to initiate a chat session with a member of the criminal group posing as a customer support representative from the exchange. The criminals will use this interface to attempt to access the users if their initial credential relay failed or if time expired, you know, because we know that these one-time passwords only are limited to 30 seconds, and then they change, so it may have expired. If so, they'll prompt the user for their username, password, and second-factor authentication code again directly in the chat window. The criminal will then take this directly to a browser on their machine and again try to access the user's account.

Should this also fail for any number of reasons, most common of which is that the device the attacker is using to access the victim's account or wallet is not, as I mentioned before, an "authorized device" in the user's profile, which probably means unknown IP or it doesn't have a persistent cookie, which the user's browser would have, even if they've said I don't want to remain logged in, they would still have a, you know, that would be a session cookie. Separately they'd have a persistent cookie which says this browser has logged in in the past. In that case, the attacker will proceed to Phase 3 with the victim. The group uses the "tawk.to" (T-A-L-K dot T-O) chat plugin on all the sites, each with the same customer support representative named Veronica. So be wary if Veronica is talking to you.

If the previous efforts have not succeeded in giving the criminal group access to the victim's wallet, they will instruct the victim to download the "TeamViewer" remote access control app. They instruct the victim that this is to help them diagnose the issue with their account directly on the user's machine. Once the victim has installed TeamViewer on their device and entered the code provided by the group, right, to initiate the session, the criminal now has full control of this poor user's device and will guide them through the steps required to authorize their device, that is, you know, their own machine, wherever they are, to the victim's account and hijack their session.

The criminal has the user navigate to their email inbox associated with the crypto exchange or wallet account. They'll instruct the user to log into their account on the exchange or wallet site. While the user is logging in, the attacker, who has control of the victim's device, will enter a random character while the victim is entering their password, right, like interject a character midstream which will force it to fail. The attacker will then click into the TeamViewer chat box without the victim's knowledge and ask them to enter their password again, which is just of course sending the password now to the criminal in plaintext.

When the user re-authenticates, the attacker will simultaneously log into the user's account on their own device, which will prompt a "new device confirmation" link to be sent to the user. The criminal then takes over the user's desktop session and sends themselves, via the TeamViewer chat feature, the device confirmation link. They can now use this link to validate their own device to access the user's account.

The final draining of the user's cryptocurrency funds may then be initiated, will be initiated during any of the previous attack phases as soon as the bad guys have access to the wallet. It's of course only contingent upon the attacker finally being able to successfully authenticate to the victim's account from their own machine being recognized as an authenticated machine, if it hasn't already been. And of course once the criminal is in the victim's account, they'll immediately begin transferring the cryptocurrency held in any of the victim's wallets to their own. And they keep the victim engaged and waiting as they steal their funds in the background on their own machine in the event that the service they're draining funds from might require some sort of email or additional phone confirmation of funds transfer. If that's the case, the attacker will assure the victim that this is normal and expected activity related to their account restoration.

Once all the funds have been sent from the victim to the criminal's wallet, they end the communication with the victim, having emptied the target's wallet. So that should give everyone a sense for how much effort bad guys in some sort of big cyber farm, you know, cryptocurrency exchange farm are willing to do to phish people who have cryptocurrency and relieve them of that burden. Amazing.

Leo: I wonder if they'll move on now that crypto's gotten less and less valuable.

Steve: I don't know.

Leo: It's nicely anonymous. It's a great thing to steal because it's hard to track.

Steve: Yes. And toward the end of the podcast I'm going to talk briefly about my own experience with having an open web server where anyone is able to create an account.

Leo: Yikes.

Steve: Leo, the Internet has become a sewer. And I know from my experience in trying to prevent that that there are - and in fact from talking to some of the anti-forum-spam people who I struck up a dialogue with, that there are rooms full of people sitting at screens and keyboards who do nothing but that all day long. And there are different rooms full of similar people who do nothing but respond to phishing cryptocurrency link clicks and then perpetrate all of this, draining people individually of their crypto currency. So, you know, it cost, I mean, if they're willing to do that to create an account against all odds on a web forum, they are certainly willing to do something not that much more in order to get a hold of someone's cryptocurrency wallet that may have a bunch of money in it. Unbelievable.

If any of our listeners are looking for something to do, the Karakurt group, with known ties to former Conti gang members, and known for its hack-and-leak extortion operations, announced this week that they are recruiting people to breach networks, code malware, socially engineer people, and extort companies for payments. And of course I'm not serious about any of our listeners wanting a job there. But their online posting was wonderful.

So just to back up a little bit, Karakurt (K-A-R-A-K-U-R-T) gets its name from a type of black widow spider. It's not a ransomware gang. They don't bother with encryption. They're known for extortion and for demanding ransoms between $25,000 and as much as $13 million, payable in Bitcoin. They don't target specific sectors or industries. They're an equal opportunity denizen. The gang backs up their claims of stolen data using screenshots and copies of exfiltrated files as proof that they've been in someone's network, and they threaten to sell or leak the data publicly if they don't receive a payment.

And they're not very patient. Karakurt typically sets a one-week deadline to pay. Until they're paid, they bully their victims by harassing their employees, business partners, and customers with emails and phone calls all aimed to pressure the company into paying the ransom. So not nice people.

Okay. Their site on the dark web is a Tor hidden service; so, you know, it's a .onion domain. It contains several terabytes' worth of previous victim data, along with press releases naming organizations that had not paid up in terms of getting ransom, and instructions for buying victims' data. The site surfaced in May. The miscreants usually break into networks by either purchasing stolen login credentials; using third-party Initial Access Brokers that we've spoken about extensively previously, you know, of course, those are brokers that sell access to compromised systems; or by abusing security weaknesses in the network's infrastructure.

Okay. So this brings us to their so-called "Great Recruitment" posting recently, last week, on the dark web. Since it was interesting and somewhat entertaining, I thought it would be worth sharing. Now, they're Russians. But I found myself thinking, wow, okay, they're not having a translation problem into English in this instance. The posting is well translated into English.

They wrote in this posting: "The Karakurt team is glad to announce some news. More than a year in private mode, but now we open the great recruitment. You can join our honorable mission to make companies pay for the existing gaps in their cybersecurity and for the inaction of their IT staff. So, our dear hack lovers, what we have for you:

"Are you an experienced pen tester and for some reason do not want to work with ransomware operators? We can find a better place in our team." Meaning they don't do ransomware. Otherwise they're every bit as evil. "Do you work for a company that you hate with all your heart? Or maybe your boss fired you, but forgot to turn off your network access? You can find solace in our arms. You are a bearer of a sacred knowledge of malware coding? Disassembling? Exploit developing? The Karakurt team is ready to set interesting and non-trivial tasks for research, implementation of specialized software, and modification of toolkits.

"Are you from the financial industry? Do you know how to make money on quotes of companies whose shares are in poor condition? Know how to sell data in a specific market? We will hug you and love you more than anyone has ever loved you before. Are you from a data recovery company and know us? Let's be friends. Maybe even best friends. Do you have social engineering experiences? There is also a vacancy. Want to take revenge on capitalism through cyberspace? We will find you both a vacancy and a psychologist. Perhaps you're a crazy researcher. We are really waiting for you, bro. The best hacker group, Karakurt, is waiting for you, our dear hack lover."

So the good news is that's not being seen by most people who are not visiting the dark web. And I assume if you're visiting the dark web you're either a security researcher who is not interested, or you're a bad guy who might be. Anyway, now you know. Karakurt has their arms wide open, ready to love you more than you've ever been loved.

Okay. And speaking of job offers, over the summer the U.S. government held what they called a Cybersecurity Apprenticeship Sprint. As a result of that, 7,000 apprentices were hired in official cybersecurity roles with around 1,000 of the new hires being sourced from the private sector. The sprint was launched in July by the White House and the Department of Labor as a way to boost the government's cybersecurity workforce.

Okay. I mentioned a web server from the Dark Ages. The security firm Recorded Future found that a Chinese Advanced Persistent Threat actor had leveraged a vulnerability in an IoT device to gain access to an electrical grid operator in India. And in a report last week, Microsoft said that they had identified the entry point for the attack. It was a tiny, somewhat obscure web server known as Boa. It's www.boa.org. And actually I was surprised that there was a three-letter dot org. Those are rare. And it's only due to the fact that it's been around for a long time. Boa, which is said to be widely used across the IoT and ICS (Industrial Control Systems) space.

Okay. As we all know, it can be very handy to have a nice simple and tight little web server, so tiny that it could even be considered a component. Although Boa is written for Unix-like operating systems, it doesn't use the traditional Unix fork and spawn approach of creating multiple instances of itself to handle individual incoming connections. I didn't study Boa long enough to determine whether it's multithreaded, thus spawning a new thread for each request. It might be purely serializing. Since the Unix/Berkeley sockets TCP/IP stack supports a queue of waiting connections, Boa might simply accept one connection after another using a single thread of execution. That would indeed make it quite lean. And apparently Boa is also quite fast. Of course you get that until you overload it by an HTTP server that is so simple.

Okay. All of that is okay. But here's the problem. It's not that Boa was first written and released 27 years ago in 1995. That's fine. The problem is that the last attention its source code received was 17 years ago, back in February of 2005. In looking through Boa's development history, I noted with some...

Leo: Is this the website?

Steve: Yes, my friend.

Leo: It looks very...

Steve: That makes mine look modern.

Leo: It's very not - last updated February 2005.

Steve: Uh-huh.

Leo: And it's, you know, I couldn't pull it up because it's not HTTPS. I had to just...

Steve: Oh, no. Nor is the web server, Leo.

Leo: Yeah.

Steve: Uh-huh. Okay. So if you click on "News," that first link there, and then if you scroll down to the 2002 Developer's Conference...

Leo: Oh, yeah, the big Boa Developer's Conference, who could forget that?

Steve: Well, yeah, in fact I have a picture of the Developer Conference attendees.

Leo: What a party.

Steve: In the show notes. I noted with some interest on...

Leo: There's just two of them.

Steve: On October 4th and 5th of 2002, the Boa Developer's Conference was held. The official minutes of the event noted: "Larry and one of his sons stayed at Jon's house October 4th and 5th, 2002. While the reasons were unrelated to Boa development, and in fact Larry and Jon spent only a few hours discussing Boa, computers, and the Free World, it seemed appropriate to refer to the event as a Developer's Conference. Here is a picture..."

Leo: Of the team.

Steve: Here is, yeah, the entire team in one location. "Here is a picture of Larry and Jon at Jon's house. (Left to right: Jon, Larry)." Now...

Leo: Oh, my goodness.

Steve: This web server is in an IoT device which is being used by the grid operators of, what was it that I said, Israel? No.

Leo: India. India.

Steve: India, right. So, you know...

Leo: Well, the price was right, I guess.

Steve: Oh, it certainly was. I have no doubt that these two have their hearts in the right place, if they're still beating.

Leo: If they're around, yeah.

Steve: But a web server they wrote 27 years ago and last tweaked 17 years ago, which has no support for secure connections, is currently in use, and apparently widely so because it's apparently very popular.

Leo: Wow.

Steve: Among other places, the operation of an electrical grid operator in India. Lord only knows where else this Boa Constrictor might be lurking.

Leo: There are a lot, I mean, you know, there are a lot of mini specialty web servers. That's a simple thing...

Steve: Yeah, it takes about an afternoon.

Leo: ...to do a lot of languages.

Steve: It takes an afternoon to write one these days.

Leo: Yeah, simple. Yeah, yeah. But wow. Why they chose this one is baffling.

Steve: Well, it's tiny. Right? So it's like, well, we're going to put it in ROM. Who's got the smallest server? Oh, look, Boa. Oh, and you didn't bring up their logo page on that site, Leo. It's pretty good. These are if you want to put a logo on your home page when you've used the Boa Constrictor server in order to serve your pages. You can pick from any of these...

Leo: I'm going to put this on my website just for fun. Powered by Boa, the high-perform - when you feel the need for speed.

Steve: Yeah, I like the one with the colored scales.

Leo: Oh, yeah, yeah.

Steve: That one's good.

Leo: That's nice. That'd look good on my site. Ooh.

Steve: Anyway, unfortunately, IoT devices on the 'Net are powered by Boa. And Microsoft didn't specify the way in, but China found a way in. And it's not surprising. I did a search on their errata page for null. And I found lots of null pointer problems in the past. So presumably not all of them. Yikes.

Leo: But good news, it's Y2K-compliant.

Steve: Yes, yes. Your concerns from 22 years ago about Y2K have been addressed. Larry and Jon did it by phone. They decided not to have a developers conference for that because - and there actually is. They go on at some length on their explanation page about Y2K. And while the underlying OS may have a problem with it, at least their code doesn't. So rest assured, if your clock is set wrong, you'll be okay.

Leo: I notice they copied their Y2K statement from the Apache Project. So I guess they were aware of that other little web server out there.

Steve: Yeah, no need to reinvent the wheel there.

Leo: No, that's right.

Steve: Unfortunately, they didn't copy their TLS support from Apache. So they don't have any.

Leo: Wow.

Steve: Wow. Okay. So the dilemma of closed-source Chinese networking products. I dislike the idea of, and I know you do, too, Leo, of banning foreign companies from selling their products to whomever wants to purchase them. And the idea that networking and surveillance cameras of Chinese origin might incorporate designed-in trojan capability, it does seem a little bit farfetched to me. Presumably, such cameras are not phoning home to China but are networked locally. So the first instant that unexplained data was caught transiting the wire, there would be hell to pay. But at the same time, we cannot prove the negative; right? We have no way of proving that there isn't any backdoor trojan capability present in Chinese network and surveillance cameras. So I suppose that the recent actions from the U.S. and the U.K. are understandable.

Last Friday, November 25th, both the U.S. and U.K. governments banned the use of Chinese networking and surveillance equipment, citing national security-related fears as the grounds for their decisions. The U.S. Federal Trade Commission has banned the import and sale of networking and video surveillance equipment from Chinese companies Dahua, Hikvision, Huawei, and ZTE. And I know that at least Dahua and Hikvision are state-owned companies. And we talked about Hikvision not long ago with regard to some badness that they were caught with.

So in the U.K., the Parliament has instructed government departments to cease the deployment of security cameras from Chinese companies on "sensitive sites" such as government buildings and military bases. British officials said the Chinese-made security cameras should not be connected to core networks, and that government departments should also consider removing and replacing existing equipment even before "scheduled upgrades."

U.S. and U.K. bans come after both countries' intelligence agencies warned against the use of equipment from Chinese companies, cautioning that Chinese equipment could be used for digital surveillance, digital sabotage, and economic espionage. Again, of course, they're not wrong. But we already do lots of even dumber things, like deploying proprietary design, closed-source voting machine technology in critical elections. You know, how do we know what those machines are doing?

Both Dahua and Hikvision had already lost a large chunk of their market in the U.S., after the U.S. Treasury department sanctioned the companies for providing the Chinese government with facial recognition and video tagging solutions in the government's efforts to oppress the Uyghurs. And I recall, as I mentioned, that Hikvision was on our radar separately for something that they were doing maybe six months ago or so.

We've talked about this a lot in the past. I noted that it was hard to believe that Russia was still using the American-made closed-source Windows OS when hostilities between the U.S. and Russia have been so aggravated. And it's also amazing that, until now, the U.S. has been deploying Chinese-made networking gear while having absolutely no idea what's inside the box. In the past we've even discussed the existence of counterfeit Cisco networking gear. Since Cisco equipment is all manufactured in China, both the real and the clearly counterfeit equipment all comes from the same place. How do we know what the counterfeit systems are going to do?

And the burden of trust is really not symmetrical. Due to Chinese massive manufacturing and fabrication capability, they receive Western technology from us, and the West purchases the resulting Chinese products from the East. Thus, more trust is required from the West than is from the East. So I suppose my point is we cannot discount such concerns as being purely hyperbolic and inflammatory. Our dependence upon our networks and digital infrastructure has slowly but surely been growing through the last several decades.

So it's only natural that at some point, someone at the national government level is going to wake up one morning and pose the big "but what if" question to their staff. You know, it's that "but what if" that was the driving factor behind the recent decision to "just say no" to Chinese networking and video equipment. And unfortunately, the protectionism that results I think is both sane and rational, even if you can't prove that anybody's doing anything wrong. You know, what if?

And, you know, the equipment we're buying is just a black box. We plug it in, and we assume it's going to be okay, but we have no ability to prove that that's the case. It really is a dilemma that we've gotten ourselves in. And all I can see is that over time, between countries where there are clear hostilities, we're just not going to be able to trust equipment from each other. And, you know, I think that's what has to happen until and unless open source ultimately wins, as I argue, and I know you agree, Leo, it ultimately should.

Leo: Oh. I didn't realize you were a complete fan.

Steve: Oh, yeah, yeah.

Leo: Good. I am, too, yeah.

Steve: Yeah. I absolutely think that's...

Leo: I think we're really learning that lesson over and over and over, frankly, yeah.

Steve: Yes, yes. MIT recently published its rankings of national cyber defense by nation. Interestingly, at the top of the list for the best defense, cyber defense, is Australia. In second place is the Netherlands. Third place goes to South Korea. And we here in the U.S., we just eke out Canada a little bit. We're in fourth place, with Canada in fifth. So those are the top five: Australia, Netherlands, South Korea, U.S., and Canada. Then the way MIT - so they did the top 20. So the way they organized it is top five is green. Then the middle 10 they lumped together. That's Poland, the U.K., France, Japan, Switzerland, Italy, China, Germany, Spain, and Saudi Arabia in descending order. And then the bottom five, they set them out separately as red, and that's in order of descending security: Mexico, India, Brazil, Turkey, and Indonesia.

So anyway, just sort of an interesting ranking. And it's interesting that Australia is solid. And they got a 7.83. This was all ranked out of 10. So they got a 7.83. The U.S. is 7.13. So a bit of a drop. Although Indonesia at the very bottom of this 20 is 3.46. So it's possible to be doing a bad job.

I just wanted to make a quick note for our listeners to be careful about Docker Hub images. It turns out that the security firm Sysdig scanned the official Docker Hub portal and identified 1,652 malicious Docker images which have been uploaded, as I said, on that official Docker Hub portal. More than a third contained cryptomining code, you know, making somebody some money, if you just run that Docker and don't pay any attention to what it's doing, while others contained hidden secret tokens that the attacker could later use as a backdoor into a server that was running a Docker and exposed publicly. Other Docker images contained proxy malware or dynamic DNS tools.

So anyway, just be careful. They are seductively easy to grab and deploy. They're very cool. But not everyone who's creating and making them available for everyone is doing so out of the goodness of their heart. So a word of warning.

We've been tracking zero-days for a while. I wanted to note that Google just fixed Chrome's eighth zero-day of the year. So they're doing better than they were last year. They updated Chrome to eliminate CVE-2022-4135 which, no surprise, was a heap buffer overflow. It was found and exploited in Chrome's GPU component. The vulnerability was discovered by one of Google's TAG researchers and is now history. So eight for Chrome, eight zero-days for 2022, and I imagine they'll get through the rest of the year. We'll see.

CISA, the Cybersecurity and Infrastructure Security Agency, is now on Mastodon, Leo. After a fake account was spotted for CISA's director, Jen Easterly, on Mastodon, CISA now has an official account on the platform. The account is at the very popular infosec.exchange server which is turning out to be where most of the industry's security researchers have been hanging out and hanging their hat. So infosec.exchange/@cisacyber is the handle, C-I-S-A-C-Y-B-E-R.

Leo: They need to add an icon and some verification.

Steve: Yeah, they didn't...

Leo: I'm not going to follow them till they put a little more effort into their account.

Steve: They didn't do very much.

Leo: Yeah. One of the nice things about Mastodon, by the way, 1,400 people already do follow them, is that it's very easy to verify that you are who you say you are. All that CISA has to do is put a Mastodon link in the CISA home page, even can be hidden, doesn't have to be visible, and they would be verified.

Steve: Oh. Very cool.

Leo: But they have so far not posted anything. They're not following anybody. They haven't put in an icon, nor are they verifying their links. But I'll take your word for it they're the real deal. You've seen this posted at CISA's site or something?

Steve: No, I picked up a news blurb about it in the infosec community.

Leo: Yeah, yeah. That is a good server, by the way. If you're in infosec, it's a good one to follow. So cisajen is not real.

Steve: Correct.

Leo: That account has been suspended.

Steve: Good.

Leo: But CISA, which is cisacyber at infosec.exchange is apparently real.

Steve: The real guys.

Leo: I'll follow them. I'll let you know if anything more happens.

Steve: And you're right.

Leo: Let's wait.

Steve: Let's hope that they go to the next step because, come on, guys.

Leo: Yeah, c'mon.

Steve: That's sloppy.

Leo: Just all you have to do is follow one person.

Steve: It's very cool.

Leo: Yeah.

Steve: Very cool.

Leo: It's good that they're there. You know, infosec.exchange has a lot of really good people on it. And I should mention that Alex Stamos, speaking of infosec, will be on TWiG tomorrow.

Steve: Oh, cool. Very cool.

Leo: Yeah. He of course was in charge of infosec at Yahoo! and then at Facebook. Left over the Cambridge Analytica scandal. Not his fault. He left because they weren't doing the right thing. And he is part of the Krebs Stamos Group. He's working with Chris Krebs now doing cybersecurity. So he'll be a great guest tomorrow on TWiG.

Steve: Yeah, Alex was first, and then they added Chris to the group.

Leo: Yeah, yeah, it's really good.

Steve: And in fact he was involved with Zoom in the early COVID-19.

Leo: That's right. He was the first person they went to when people got mad at them for not doing it right, doing encryption right, or kind of misrepresenting their encryption. He's also a professor at Stanford. So I think he will be a good guest.

Steve: Neat. Tomorrow.

Leo: Yeah, yeah.

Steve: I have one piece of miscellany. It's not directly security-related, or privacy. But everyone's talking about Twitter and its uncertain future under the reign of Elon. I stumbled upon something that I thought our listeners might find interesting, and I think you might, Leo, as I did, because it appears to contain some actual facts. This is a note written by an unnamed executive director at an unnamed business-to-business organization. But it looks authentic. I presume it's anonymous because he would prefer not to have Elon Musk retaliate against his firm. The title of his posting was "I told my team to pause our $750,000 per month, so three quarters of a million dollar per month, Twitter ads budget last week."

So here's what he wrote. He said: "I've seen a lot of technical and ideological takes on Elon Twitter." And I got a kick out of that. I wondered whether it was a play on "Tim Apple." Anyway, he said: "But I wanted to share the marketing perspective. For background, I'm a director at a medium-size B2B tech company, not in financial services anymore, running a team that deploys about $80 million in ad spend per year. Twitter was 8-10% of our media mix, and we have run cost-per-engagement, i.e. download a white paper, register for an event, et cetera, campaigns successfully since 2016.

"I had my team keep our Twitter campaigns live for two weeks post-takeover on the bet that efficiency would improve with fewer advertisers, and that the risks were managed and probably overblown. I was wrong, and I think the things we saw in these last two weeks means many more advertisers will bail on the platform in the coming weeks," and he says, "(for non-ideological or virtue-signaling reasons)."

So then he has four bullet points. He says: "Performance fell significantly. CPMs didn't drop" - meaning same number of eyeballs - "but our engagement went way down. Maybe it's a shift in users on the platform. Maybe it's ad serving-related."

Second point: "Serious brand safety issues." He said: "Our organic social and CS teams got dozens of screenshots of our ads next to awful content. Replies to our posts with hardcore anti-Semitism and adult spam remained up for days even after being flagged." Third: "Our entire account team at Twitter turned over multiple times in two weeks. We had multiple people," he said, "AE, AM, analyst, creative specialist, supporting our account, and they all vanished without so much as an email. We finally got an email with a name for an AM" - I guess that means account manager - "last week, but they quit, and we don't have a new one yet."

And, finally, he said: "Ads UI is very buggy, and login with single-sign-on and two-factor authentication broken. One of my campaign managers logged in last week and found all our paused creatives from the past six years had been reactivated. Campaign changes don't save. These things cost us real money." Anyway, I thought that was...

Leo: I wonder if they put any prices with the decimal point in the wrong place up by accident. Now, that could cost you.

Steve: Since I hadn't encountered anything as substantive as that, I thought that it was interesting to see. And I understand a bit about what's going on from the perspective of one of Twitter's advertisers who views the service dispassionately. He doesn't care one way or another who's doing what, except he dislikes the idea of their ads appearing to endorse horrific content which it's now appearing next to or in the comments that follow an ad. You know, for him, Twitter is just either an ends to a means - wait. A means to an end. Or maybe not. So anyway, I thought that was interesting.

Leo: Yeah, he's just a businessperson; right.

Steve: Yeah. Oh. And in a related piece in a security newsletter I recently scanned, the statement was made: "Some threat intelligence companies are telling their customers that they can no longer guarantee takedowns of malicious or reputation-damaging content from Twitter as there is nobody in Twitter's abuse team to respond to requests anymore." So another data point from a different direction. And for what it's worth, TweetDeck is behaving weirdly now. You know, I always go in in order to pull feedback from largely my DMs, although I scan the public feed, you know, the @SGgrc postings. And it was definitely not working the way it used to, and not in a way that I liked. So something is changing or has changed. And, you know, I don't know, I don't care to know what that is.

Okay. So KerryOnAnon is his name. MrIndigo is his Twitter handle. He said: "Hi, Steve. Finally listening to the latest Episode 898, and I started wondering, is quantum computing going to be just a faster way to guess passwords, or is there another attack vector? In other words, is it just going to be a faster way to brute force attack passwords?"

Okay. Interestingly enough, once we get quantum computing, assuming that we ever get quantum computing, it won't be any faster at brute forcing passwords. In fact, it would likely be far slower and vastly more expensive than conventional hardware-accelerated, hash-based password brute forcing.

Leo: Oh, how interesting. That's not the problem.

Steve: No. There's just a class of things it's good at. The rest it's really crappy at. You know, it's like weather prediction. It can do that, but it can't tell you where a specific drop of rain is going to land, and that's what you need for symmetric crypto and hashing is that kind of exact operation. The important thing to understand here is that some of today's crypto, but only some of it, depends upon the traditional, time-proven difficulty of factoring a very large number into its two half-as-large prime number components.

That's it. That's all that the fervor surrounding quantum computing is about, the ability to do a couple of things quickly that are entirely insurmountable, that is, this factorization problem. But it's only the asymmetric key crypto that quantum computing might be able to someday weaken. None of the other crypto that we also depend upon today will be affected. Symmetric key crypto, like our beloved AES ciphers or today's strong hashing algorithms, will not be affected at all. And they don't need to be changed.

I was thinking about quantum computing after I read this guy's note, and I was looking for a good analogy of the effort, its promise and the difficulty that it presents. And what popped into my head as being in almost every way similar was power generation, at scale, via nuclear fusion. It's a useful analogy. It requires crazy, way out there, new physics and new materials and new technologies. And like quantum computing, fusion has been chased for decades, driven by the promise of "what if," just like quantum computing has. And incredible amounts of ingenuity and money have been sunk into it. Many different approaches have been tried and discarded.

And yes, we are creeping forward little by little, inch by inch, tantalizingly, just enough to keep the investment cash flowing. But, boy, is fusion a difficult nut to crack. In order to fuse matter, we must create, contain, and compress the hottest plasmas humans have ever handled - hotter, it turns out, than the sun. And at this point it's as much art as science. Will we get there someday? Maybe. Maybe not. It's still not clear. But as with quantum computing, we do appear to be making some progress year after year, learning as we go.

So as for quantum computing, my feeling is that there's no reason not to replace that small but crucial portion of our large crypto library of algorithms, which are believed to be currently unsafe if quantum computing ever happens. We can replace it with algorithms which are believed to be quantum safe. We just don't want to make any mistakes with our replacements, and there's no reason to believe that there's any big hurry. We might well have free electricity, once we figure out how to burn water, before quantum computers threaten our current dependence on today's asymmetric crypto. So not to worry.

Another listener who requested anonymity, and I'll explain why in a second, he said: "Hi, Steve. In the last episode of Security Now! you talked about passkeys.directory, which lists web applications that support Passkeys. I wanted to share my observations with you. First, the website owner chose to manage it with no transparency. When I saw it, I thought there must be a Git repo where I could open an issue for a change request. Surprisingly, they choose to use Google Forms, which masks all the review and approval process." And he's talking about passkeys.directory.

"Second," he said, "I've noticed that many companies in this list are also customers of OwnID, which is listed as the authentication provider, including Carnival Cruises."

Leo: Oh, interesting.

Steve: Yes. Yes. They did not do it natively. And he says: "And then investigating the OwnID flow." He said: "When Leo pressed the fingerprint button, the QR code encoded a URL that sent his iPhone to passwordless.carnival.com with a session identifier. Then he performed a WebAuthn authentication on his iPhone. Once completed, the session got updated on the server, and the browser on his laptop logged in. The flow is using WebAuthn's Passkeys, but not likely the way it was designed to be used. WebAuthn phishing resistance mechanism works in a way that a Javascript API called on the browser triggers the underlying library and matches the domain a key was registered in and the domain asking to authenticate.

"By implementing WebAuthn as it is in Carnival, the phishing resistance mechanism suffers from a flaw. As an attacker, you can spoof Carnival's login page. So the user sees the same page, only a different domain. When you click the Biometrics button, the attacker's backend will send a request to Carnival to get a QR code which encodes the passwordless.carnival.com. Then the phone would ask you for your face or fingerprint to authenticate with a Passkey, which will update the session on the backend, and the attacker gets in. Actually this is a thing that I spent a lot of time on SQRL solving completely, and it's crucial."

He says: "The right way to implement Passkeys is by calling the WebAuthn API on the laptop's browser," he says, "instead of presenting the QR that will open a browser on the mobile phone, and letting the browser do its job, presenting native WebAuthn screens, including a QR which is scannable from a mobile phone. This way, the domain you're authenticating to is passed in a side channel" - that is, you know, Push versus BLE, Blue Tooth Low Energy, from the browser to the phone - "to the mobile phone directly from the browser, and a phishing site will be blocked as the credential on the phone was registered under the original domain."

Okay. So first of all, our listener who wrote this to me is 100% correct. And by the way, he's a developer for an authentication provider who asked for anonymity. Another way to say this is that, rather than doing the work of upgrading their own servers to become a first-party Passkeys provider, Carnival Cruises, and unfortunately a lot on that list, has outsourced their authentication responsibility to a third-party provider, in this case OwnID. But in doing so, by punting in this way, they've bypassed Passkeys phishing protections. This gives their visitors the false belief that they're getting the hack-proof benefits of Passkeys without actually getting them. This could be transient. We can hope not.

But on the other hand, OwnID is in the business of doing this. So they're going to presumably keep selling their "instant onboarding" services, and most websites will simply want easy login without really caring about their visitors' security. So we've seen the first way that Passkeys will fail. And that is, when implemented like this, you can be phished. And that was a big deal. It was supposed to be anti-phishing. Well, it's only anti-phishing if you don't turn the responsibility over to a third party. And if you do, and this page of people have, you're not getting the benefit of Passkeys.

Leo: Oh, that's disappointing.

Steve: All you're getting is - yeah.

Leo: But of course to be predicted.

Steve: Yeah, exactly.

Leo: Yeah.

Steve: Christopher Ursich, he said: "SN topic request: Hardware Security Modules." He said: "You said you had one. Besides the technical crypto, can you describe how you interact with it in practice to sign your code?"

Sure. Just as there are EV (Extended Validation) TLS certificates for web servers, there are EV code-signing certificates. I have no idea whether they are any better or more trusted than non-EV code-signing certificates. But I'll take every advantage I can get. And one requirement of EV code signing is that they must, without exception, be protected by a hardware security module so that the EV private key can only ever be used for signing, and cannot possibly escape into the wild.

The EV code-signing key which I purchased from DigiCert was packaged in a Gemalto USB dongle which is paired with the SafeNet Authentication Client. Somehow, when I use the same Authenticode code-signing command in Windows as I've always used, that SafeNet client is invoked, the hash of the file I'm signing is sent to the key and signed inside there, and it returns a signed blob. So it's just a matter of having a free USB port and installing a hardware interface client.

Part of the effort which I'll be engaged in toward the end of the work to publish the final SpinRite 6.1 code, which will be like 6.0 is, a hybrid DOS and Windows app, will be automating this code-signing process server side. Since each owner's copy of SpinRite embeds their license information, which makes their executable unique, each one needs to be individually code-signed on the fly by the server as it's downloaded.

What's going to be really annoying is that Windows Defender will always be complaining, for every single user, that the user-specific custom SpinRite file is not commonly downloaded, thus needlessly warning and alarming its users. We've seen that no degree of reputable signing is able to bypass this alarm. I discovered that when I signed the final version of SQRL, and I updated the DNS Benchmark. You know, people said, "Hey, Windows Defender's not happy." And I said, "I know." Doesn't matter if you sign, and those were EV certificate signed. Windows Defender says, "Oh, haven't seen this a lot before."

And you could understand, it's going to take a hash of the things that you want to download, and it's obviously sharing those in the cloud. And when it sees enough of those and no complaints, then it goes, okay, it must be okay, and stops bringing up warning messages. Unfortunately, SpinRite's users are just going to have to get used to that because every one of those that they download is going to be unique.

Two people. Dangard asked: "Steve, how can I get access to test the pre-release version of SpinRite 6.1? Feel free to email me or just respond here. Thanks so much for your work on SpinRite. I have drives waiting for 6.1."

And sdholden asked: "Hey, Steve. Not sure the best way to reach you about the Git server for SpinRite, so I thought I'd start here. When I try to create an account, I get a dialog box asking me to sign in instead of allowing me to create a registration," he says, dot dot dot question mark.

Okay. To both listeners and everyone else, in case some of you hadn't noticed, the Internet has, sadly, become a sewer full of both bots trolling constantly and even human labor farms being paid for creating accounts online. I've been running two web forum servers for years. Despite having all manner of entrance barriers erected, like even requiring the correct answer to the question "What software is Steve best known for?" in order to create an account, five out of six of the account registrations were bogus in those forums. Like, how does a bot know?

Leo: How hard is that? How hard is that? Yeah, a bot wouldn't know, but...

Steve: No. I know. At one point we had 6,500 users registered in GRC's forums. And I was thinking, wow, I haven't even talked about it that much. Okay. Now that number is a bit over 1,100, after I spent several days working to get that under control.

Leo: Yeah.

Steve: 5,500 of those were registered in Afghanistan and Turkey and Indonesia. I mean, just like - and Russia. And, you know, just it was so annoying.

Leo: Spammers love forums. They really do.

Steve: Oh, my god, yes. So I've erected much tougher barriers since, and I've mostly gotten it under control. And since I erected those stronger barriers, 20,204 additional account creation attempts have been thwarted. So I have an additional 20,000 bogus users on top of the 5,500 I had before. The reality is that, today, as you said, Leo, running any sort of open web service results in a torrent of bogus registrations. And even with all that in place, the wonderful volunteer moderators I have, who make time to read everything, are still removing users who attempt to subtly pollute our content.

So here's the problem. GRC's forums need to be open. So I have no choice other than to erect the strongest account creation barriers I can, then apologize to those whom we mistakenly reject as false positives, and also weed out those who do slip past the barriers due to false negatives. But GRC's GitLab server has no need to be open. So it's closed. Its account creation page is protected by a magic incantation which must be provided before the troll that guards the bridge will allow newcomers to pass. It requires insider information which can only be obtained by participating in GRC's old-school, blessedly wonderful text-only NNTP newsgroups. Once someone shows up there and is able to post, they can ask how to satisfy our cantankerous GitLab troll. But also note that we're not using GitLab for any social interaction. We're only using it for issue management.

At this point, what I need is feedback from people who are testing SpinRite 6.1. Since we have a handful of known issues to fix - I'll get to that in a moment - it's best for newcomers to join and catch up on all the various threads in the newsgroup in order to eliminate duplicate postings of already-known problems. So if anyone is really and truly interested in participating in SpinRite 6.1's testing, you're invited to head over to GRC's Discussions page - that's the page at GRC, if you google "GRC.com discussions" it'll take you there - and create a connection to our news server. Find the grc.spinrite.dev group and say hi.

And speaking of SpinRite, it's working. As I planned, I updated GRC's primary server to handle downloading of pre-release versions of SpinRite, and last Friday morning, after Thanksgiving, I posted the information in GRC's spinrite.dev newsgroup about where any existing SpinRite owner could go to grab their own copy. I'll share three newsgroup anecdotes which I've edited just a bit for podcast clarity.

A few hours after my first release announcement, someone whose handle is DarkWinX posted on Friday at 2:44 p.m.: "Well I can already report success with a USB. In my race to find something to eagerly test on, with the short time I had, I grabbed an old USB I received with the purchase of StarCraft II. I figured I'd reformat it with InitDisk and run SpinRite from there. So I put it in the computer and started InitDisk. It waited, and waited, for about 30 seconds. Eventually the USB was recognized by Windows and showed up so I could nuke it. I tried it again and it still took around 30 seconds to load. So I figured maybe not the best USB to run SpinRite from. So I found another. I thought, why not run SpinRite on the problem USB as a target, so that's what I did. After a Level 2 scan, without finding anything wrong, I rebooted, plugged it in, and instant success. That USB now loads inside Windows instantly every time. Looking forward to testing some more."

Second comment, Saturday morning, 8:39, Mark Ping posted: "Finished the Level 2 in two hours for a 1TB. Then ran Level 4, and it took nine hours, 37 minutes for 1TB, compared with 150 hours before." And then he finished: "SpinRite is back, baby."

And finally, Dale F., Saturday evening at 10:12 p.m., posted: "I have a 500MB laptop drive that I put in a SABRENT portable enclosure. After I dropped it about two years ago, it could not be recognized by any PC, or by SpinRite 6.0. So I said to myself, 'Just have to wait for 6.1.' On Friday, I ran a Level 2 with SpinRite's first alpha release, and one hour later it was good as new. Thanks, Steve."

Okay. So frankly, SpinRite's first functional pre-release debut could not have gone much better, and it went far better than it might have. Over the weekend, using the feedback provided by the large group of avid testers, we moved SpinRite through three more releases to its fourth alpha release by mid-afternoon on Sunday. And with only a few exceptions, it is now working well for everyone. Overall, it's 100% functional in every way that matters. There are a number of things that I need to fix, like SpinRite's various clocks are not continuing to operate while it's deep into data recovery. I recently rewrote that entire data recovery system, and I just forgot to periodically update the clocks while I was in there. So actually I'm going to change the entire way that works so that it's much better.

Another example is that SpinRite's predictions of its remaining time to run is not working right when it's started midway into a drive rather than at the beginning. You could start it wherever you want to. Anyway, it was working once, and something I did broke that. So I'll fix that. So right now the newsgroup gang is continuing to pound away on the fourth alpha release, logging everything they encounter in our GitLab instance. While that's underway, my own now highest priority is to make a decision about that next operating system that I'm considering purchasing and moving to. Its licensing deadline, as I mentioned before, is the end of the year. It's either by then or never.

So I expect that to take - that's what I'm going to be doing this evening. I'll start that. I only think it'll take a couple days. I just want to make sure that I can boot something, you know, the classic "Hello World" app, both from a BIOS and from a UEFI-based machine. Then that says yes, I'm going to go with this OS. Then I'll return and get SpinRite's DOS executable completely finished.

I should mention, I told you this, Leo, before we began recording today. One thing happened this morning that completely caught me off guard. I hired Greg, who everyone has heard me refer to through the years, 32 years ago tomorrow. Tomorrow is his 32-year anniversary of employment with GRC. That means that tomorrow he will have been providing technical support for SpinRite for 32 years. Yesterday he fired up the latest SpinRite 6.1 alpha. He'd never seen it before. He's seen nothing until, you know, I've been keeping him and Sue appraised of what was going on. I sent them both an email saying, "Well, it works, somewhat to my amazement."

So he fired up the latest SpinRite 6.1 alpha, ran it on a bunch of drives he had around. He said that he ran it on a 1TB spinner which took about two hours. That's about right. Remember I've thought about half a terabyte per hour is good performance for a spinning drive. And that certainly beats two weeks. And still it wasn't instantaneous because it was a spinning drive. Then, he said, he scanned a 128GB SSD in five minutes, and he was stunned.

So he told me on the phone this morning that he knows the number one question he is certain people are going to be asking, once SpinRite's previous users start using 6.1, is how SpinRite 6.1 could possibly be so much faster? It was like, it's like the difference is too much to believe. You know, either 6 was like way slow, or is 6.1 actually doing anything? On the other hand, I should also mention that a number, a whole bunch of people in the newsgroup have actually had it recovering data, recovering drives, things that could never be copied before. We're seeing green R's on the map showing data was problematical and was recovered. So anyway, I'm very excited that I will be able to soon stop talking about it and have it in everybody's hands.

Leo: Woohoo.

Steve: Yeah.

Leo: Very, very good news. Thank you for the hard work.

Steve: Well, thank everybody for their support. I really appreciate it.

Leo: Yeah. So you said NNTP, your newsgroups are NNTP. I thought it was XenForo. Or does XenForo use NNTP? Is that why?

Steve: No, XenForo is the web forums.

Leo: Oh, you have newsgroups in addition to the web forums. I get it.

Steve: Yes.

Leo: I get it.

Steve: Yes. Newsgroups I've had forever, and I love them. They're little backwater...

Leo: Yeah, they sure are.

Steve: We get real serious work done.

Leo: How do you read a newsgroup these days?

Steve: Thunderbird is a really good newsgroup reader.

Leo: Okay.

Steve: It does a good job of it. On the discussions page, I list - I asked the question of everybody, like six months ago, and there's like a list of maybe 30 different NNTP clients. There's only one for iOS which is called NewsTap. It's a great little newsreader for iOS. There's a bunch of newsreaders for Android, and a bunch of Linux and Mac and Windows.

Leo: And then you host it. It's on your GRC site.

Steve: Right. It's news.grc.com.

Leo: Nice.

Steve: And that's been one of the things I've had, well, okay. So here's the reality. SpinRite 6.1 will ship. It will be perfect. The newsgroups are why it will be perfect. In this day and age, once upon a time, you know, back when we had DOS 2 or 3 or 4, I could write a program, and it would work everywhere. Those days are gone.

Leo: Yes.

Steve: I could never...

Leo: You need testers nowadays

Steve: I could never do this if it weren't for the guys in the newsgroup. As I said before, I've got, like, all these motherboards around now, and all these old hard drives, because it was like, uh, Steve, the ASUS Cranox 327 isn't working. So I'd go onto eBay, ASUS Cranox 327. Oh, yeah, there it is. And I buy it. You know? So Lorrie is saying, do we still need all these? I go, no, honey, just a little bit longer.

Leo: A little bit longer.

Steve: A little bit longer.

Leo: Little bit longer. Yeah, used to be that all the browsers could handle newsgroups. But they've slowly stripped that out of every browser.

Steve: And of course FTP is gone now, too.

Leo: It's gone, too, that's right. They take all - and reasonably so. If nobody uses it, it's just a security...

Steve: But a good generic newsreader is Thunderbird. It's multiplatform, and it's pretty good for getting the job done.

Leo: Good. I'll have to check out the newsgroups. For some reason I spaced that you have a newsgroup. I thought it was all forums. Which forums are fairly old-fashioned. Newsgroups are positively antediluvian. That's good. I like it.

Steve: Yes. And the forums are where support will be for SpinRite.

Leo: Right.

Steve: I'm going to engage community support. But I'm never going to allow, I mean, like the newsgroups are my sanctum sanctorum. Is that the right term?

Leo: Do you still - does you use UUCP and send it off and everybody in the world gets to see it? Or is it just hosted on your site?

Steve: Actually, we block it going anywhere else.

Leo: Yeah. Yeah, yeah, okay.

Steve: Because Google Groups would like to be pulling from an NNTP server. The problem is people were responding to postings that Google had sucked out, and nobody was ever seeing their responses.

Leo: Right, right.

Steve: So it is closed. I actually have a technology where the IP address of the entity which pulls the article is added to the headers. So if we ever see postings out in public, we can look at the headers and see the IP address that is pulling them, and then I block them.

Leo: Oh, so smart. So there. Wow. So it's really, I mean, to call it a newsgroup is really not exactly right because the whole idea was newsgroups were federated, and they would be copied every night from university to university.

Steve: Oh, and we've got - I've written a whole bunch of extra code.

Leo: You just use the NNTP protocol for your server.

Steve: Yes. We have something called a CECIL-ID, which is also added to a posting, which is a hash of the person's username and password, which allows the postings to be owned by them. Nobody else can delete them, but they can delete their own.

Leo: Perfect.

Steve: And there's a whole bunch of other, you know, benefits that we've added over time.

Leo: Yeah, yeah. Very interesting.

Steve: I just, you know, I will - that's what I'll be using, like...

Leo: Forever.

Steve: When somebody comes along to turn off the servers after I'm gone, they'll be shutting down the newsgroups.

Leo: Oh, that'll be sad. All right, Steve. Always a pleasure. He does it the old-fashioned way. He does it the old way. But the old ways are often still the best. Steve Gibson is at GRC.com, along with his newsgroups. That is the Gibson Research Corporation. You'll find SpinRite there, the world's best mass storage recovery and maintenance utility, now faster than ever. It's really working. It is. It's really doing something, honest. If you don't have a copy, get 6.0 now, you'll have a free upgrade to 6.1 when it comes out. You can also participate in the development and all of that, as he said. GRC.com.

While you're there, you can get a copy of this show. Security Now! is hosted at TWiT.tv but also at GRC.com. Steve has two unique versions, a 16Kb audio version for the bandwidth-impaired. He's always done that from day one. And for his transcriptionist, actually, Elaine Farris, because she writes this all out, and she's living in the country with a lot of horses, doesn't have a lot of bandwidth. You can get the transcripts there, as well, GRC.com, as a 64Kb audio file. We have audio and video at our website, TWiT.tv/sn.

There is a YouTube channel for Security Now!. That's a great way to introduce somebody to it, or if you hear something on here you want to share with other IT professionals, your boss or friends, your spouse, then just clip it at YouTube. That's probably the easiest way to do it. They make that a fairly simple thing to do. Of course subscribing in your podcast client might even be the best way to get it. That way you'll get it automatically the minute it's available. You can build your collection of all 899 episodes. Whew. That's a lot of episodes.

Steve, we will be back here next Tuesday, 1:30 Pacific, 4:30 Eastern, 21:30 UTC. Had to do the math. You can watch us live, live.twit.tv. Chat with us live at irc.twit.tv. Or if you are fortunate enough to be in the Club, you can do it in the Club TWiT Discord. Actually, you should join the Club, if you're not already a member. It supports Steve's efforts plus everything we do here. $7 a month for ad-free versions of the show, access to the Discord. You also get stuff that we don't put out in public, like Hands-On Macintosh, Hands-On Windows, the Untitled Linux Show and all of that. Thank you, my friend.

Steve: And Leo?

Leo: Yes.

Steve: Happy Birthday again.

Leo: Thank you.

Steve: For your 66th. I want you to hold onto that sign so that in 33 years you can turn it upside down and celebrate 99.

Leo: 99 etouR. Good thinking, Steve. I'll save that. I bet you save old calendars, too, don't you.

Steve: No.

Leo: Thank you, Steve. Have a great week. We'll see you next time on Security Now!.

Steve: Bye.

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/