SECURITY NOW!

Transcript of Episode #896

Something for Everyone

Description: This pure news week we look at Dropbox's handling of a minor breach, and we follow up on last week's OpenSSL flaws. The FTC has had it with a repeat offender, and we know how much total (reported) ransom was paid last year. Akamai reports on phishing kits, we have some stats about what Initial Access Brokers charge, and we look at the mechanics of cyber bank heists. Several more DeFi platforms defy belief, Russia is forced to move to Linux, the Red Cross wants a "please don't attack us" cyber seal, nutty Floridians get themselves indicted for a bold tax fraud scheme, is China cheating with zero-days, the NCSC will be scanning its citizenry, and more.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-896.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-896-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We'll talk about how Dropbox properly handled a minor breach, and ask the question of whether you should ever trust a managed service provider. More on the OpenSSL flaws. The FTC going at it with Chegg. I'm glad to see this. And is China cheating with zero-days? That and a whole lot more coming up next on Security Now!. Stay tuned.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 896, recorded Tuesday, November 8th, 2022: Something for Everyone.

It's time for Security Now!, the show where we cover you, your privacy, your security, how the Internet works, how computers work, with this guy, this genius right here, Mr. Steve Gibson. Hello, Steve.

Steve Gibson: Yo, Leo. Great to be with you. We have - this is a Patch Tuesday. This is the - it hasn't fallen on an Election Day since 2016. Just a little bit of trivia there...

Leo: Fascinating, yes.

Steve: ...for those who are following along. And we owe today's show title to my wife. Lorrie and I were out walking yesterday, and I was telling her what progress I had made so far. She said, "So do you have a topic?" And I said, "You know, I don't so far." I said, "But that's okay. Sometimes nothing really jumps out or stands out or needs special attention, and so I just call it like a busy news week or something." And she said, "How about calling it 'Something for Everyone'?" And I said, "I like that."

Leo: I like it.

Steve: And so that's today's title, "Something for Everyone." Because we just have all kinds of stuff. We've got one of our pure news weeks. We've got Dropbox's handling of a minor breach. We follow up on last week's OpenSSL flaws. The FCC has had it with a repeat offender. And we're going to find out how much total reported ransom was paid last year to the ransomware denizens. Akamai has reported on phishing kits, and that's some - it's, like, frightening. We've got some stats about what Initial Access Brokers charge. And we look at the mechanics of cyber bank heists, like how that's actually pulled off in the real world. We've got several more DeFi platforms defying belief. Russia is forced to move to Linux, finally. The Red Cross wants a "Please don't attack us" cyber seal. We've got nutty Floridians who have gotten themselves indicted in a bold tax fraud scheme that you just can't imagine they could have possibly thought they could have gotten away with. And, well, because of indictments they didn't.

Leo: You know how that is, yeah.

Steve: That's right. Also the question has been raised by Microsoft whether China is cheating with zero-days. And in what I think is a fabulous idea that I hope the U.S. might adopt, the NCSC will be scanning the U.K.'s citizenry for vulnerabilities and working with them to remediate them. And that's not all. There's more. We've got a great Picture of the Week. I've got some feedback from our listeners and a brief update on where SpinRite stands.

Leo: Oh, wow.

Steve: So as I said, something for everyone.

Leo: Something for everyone. That sounds like an excellent show. I'm looking forward to it. I always do. Do we have a picture? I didn't even look.

Steve: We have a wonderful picture. I will lead up while you're getting it ready. Now, I'm tempted to call this the dumbest thing I've ever seen except that we've got two previous occupants for that slot. One is the locked gate standing alone out in the middle of a meadow with a path running up to it. And it's like, what is this locked gate doing out in the middle of nowhere? Who's not going to walk around it? And sure enough, there's like a dirt-trodden path on either side.

The other dumbest thing was that generator that had to be grounded, so someone stuck a piece of rebar into a pail of dirt and hooked the ground wire to the rebar. And it's like, okay, I don't think that's quite what they had in mind when they said you need to ground this generator.

Okay, here we've got a very tall gate which looks like it's an electric gate.

Leo: It's a good-looking gate. It's very nice.

Steve: Nice-looking gate. Got an intercom on the side so you can buzz the person, looks like maybe three different units are back there somewhere. And you are not supposed to get in or out, presumably. Unh-unh. No. The problem is that the genius who designed this gate used a series of horizontal bars. And so I gave this the caption "Can't get in? Hmm. How about use the built-in ladder?" Because, I mean, it's like designed for scaling the gate. It's just, you know, hmm, I can't get in. What should I do? Oh, look, it's a ladder.

Leo: How handy. How convenient.

Steve: I mean, all they had to do was make them vertical, and then you'd just be like stuck. You'd be looking like, you know, like prison bars. But no, they built a ladder from the gate, and so it's quite easy. Though this goes down, this is maybe the third dumbest thing that we've seen on the podcast where the...

Leo: It's in the list, definitely.

Steve: Yeah, we are acquiring them over time. Okay. So last Tuesday, which was the first of November, Dropbox posted of their own experience titled "How We Handled a Recent Phishing Incident That Targeted Dropbox." And the short version is I think they handled it pretty well. But there are some lessons to be had surrounding the event. Their announcement began with sort of the required "do not worry" disclaimer. They said: "We were recently the target of a phishing campaign that successfully accessed some of the code we store in GitHub. No one's content, passwords, or payment information was accessed, and the issue was quickly resolved. Our core apps and infrastructure were also unaffected, as access to this code is even more limited and strictly controlled. We believe the risk to customers is minimal. Because we take our commitment to security, privacy, and transparency seriously, we've notified those affected and are sharing more here."

Okay. Then I skipped over a bunch of background. And the part I wanted to share with our listeners was this. They said: "At Dropbox, we use GitHub to host our public repositories as well as some of our private repositories. We also used CircleCI for select internal deployments." CI is some automation technology, CI standing for Continuous Integration. So they said: "In early October, multiple Dropboxers received phishing emails impersonating CircleCI, with the intent of targeting our GitHub accounts. A person can use their GitHub credentials," they explained, "to log into CircleCI."

They said: "While our systems automatically quarantined some of these emails" - you know, phishing emails, right - "others landed in Dropboxers' inboxes. These legitimate-looking emails directed employees to visit a fake CircleCI login page, enter their GitHub username and password, and then use their hardware authentication key to pass a One Time Password to the malicious site." And as we know, all of this bypasses, you know, I mean, this approach will get around the use of one-time password authenticators. So they said: "This eventually succeeded, giving the threat actor access to one of our GitHub organizations, where they proceeded to copy 130 of our code repositories." Whoops.

They said: "These repositories included our own copies of third-party libraries slightly modified for use by Dropbox, internal prototypes, and some tools and configuration files used by the security team. Importantly, they did not include code for our core apps or infrastructure. Access to those repositories is even more limited and strictly controlled." And finally: "On the same day," they said, "we were informed of the suspicious activity." They don't indicate how, but this is why you need to do network monitoring like, Leo, you were just talking about with that previous sponsor.

They said: "The threat actor's access to GitHub was disabled. Our security teams took immediate action to coordinate the rotation of all exposed developer credentials, and determine what customer data, if any, was accessed or stolen. We also reviewed our logs and found no evidence of successful abuse. To be sure, we hired outside forensic experts to verify our findings, and reported this event to the appropriate regulators and law enforcement."

Okay. So there are three points that I wanted to highlight from this report. The first is that we have yet another instance of a major security-savvy and network-savvy organization - you know, Dropbox; right? I mean, they know their way around or they wouldn't still be around - being successfully attacked and breached, even in the face of knowing that this is going on. Their email filters worked to prevent their employees from being subjected to this error-prone event mostly. But those filters also failed just enough to allow bogus phishing attacks to reach their employees.

And notice that these were code developing employees, you know, not, for example, less sophisticated clerical or office workers who you might have in a huge organization that wouldn't be expected to be up to speed on computers. You know, these are people who like log into CircleCI and GitHub, and they were fooled. The point is phishing. And we'll be talking about that several more times before the end of today's podcast.

The second point I want to make is the introduction of a new concept which I would term "the phishing email attack surface." We're all familiar with the traditional concept of an attack surface; right? The idea being that the more potential points of entry that exist, the greater the threat that any one of those might be inadvertently left open or somehow breachable. So this new concept that I would call the phishing email attack surface uses this recent Dropbox experience as a perfect example, noticing that the more complex an organization's setup is, which is to say the greater number of ancillary services an organization employs, the greater is their phishing email attack surface. There're just more things that have logons and authentication requirements and, again, more points of entry.

The modern trend is products as managed services, where companies are increasingly contracting out for an increasing number of services, rather than rolling their own in-house. The theory of this is sound. Why reinvent the same wheel over and over, especially when there's little additional value to be added by doing so? Just contract for this or that service while focusing upon the company's core mission, rather than wasting time on developing and running all of those other things that are common to all companies. Sounds great.

But recall all of the downstream damage that the breach at SolarWinds created. SolarWinds was a provider of exactly this sort of outsourced services model. And also remember all of those dental offices that were being breached, and the hospital services that were hit by crippling ransomware when their MSP, their managed service provider was breached. The danger represented by managed service providers is exactly what I'm referring to here.

So I wanted to observe that we, as an industry, still have a serious problem with remote network services authentication. The very fact that phishing emails even exists as a security issue demonstrates that this serious problem has not yet been solved. So the more remote network MSP services an organization maintains, the greater their phishing email attack surface will be.

The third and final point I wanted to make was where Dropbox wrote, they said: "On the same day we were informed of the suspicious activity, the threat actor's access to GitHub was disabled. Our security teams took immediate action to coordinate the rotation of all exposed developer credentials and determine what customer data, if any, was accessed or stolen. We also reviewed our logs and found no evidence of successful abuse."

To that I say bravo. When we were all growing up, our elementary schools conducted periodic fire drills. Without warning, alarms would sound throughout the school, and the entire school, class by class, would file out in an organized manner to previously designated locations. While I was in school, those alarms never went off except for during drills. But if someday they were to, the entire school was prepared.

My point is every organization must now be prepared for the possibility of a network breach. So "breach drills" should become a thing that all responsible organizations conduct, just as fire drills were once that when we were in elementary school. Just as when a school might be on fire, after a network intrusion we've seen the stats showing that time really can be of the essence. So planning for a breach, including having some drills, should be something that responsible organizations do. Dropbox's immediate response showed that they were ready and prepared for that eventuality. And again, I think that is of crucial importance.

Leo: I think it's also important to point out that that's probably why in many cases it's better to use an MSP than do it on your own. I mean, if we were to count all the flaws that people introduce themselves by trying to do it themselves, that's going to far outweigh the number of exploits because MSP was taken advantage of; right?

Steve: I mean, I think that's a useful consideration. The problem with an MSP is the single point of failure. So a breach at SolarWinds...

Leo: Gets everybody. Yeah, yeah, yeah.

Steve: ...devastated, yes, so many clients.

Leo: But I think about Bitwarden, for instance. And some people, again, with Bitwarden, one of our sponsors and a password manager, host their own. And they often say, well, why do you let Bitwarden host it? Because, I always say because I think they're going to - yeah, I could host it myself. I think they're more likely to keep it locked down than I am. You know?

Steve: And backed up.

Leo: Yeah.

Steve: You don't risk losing the cloud presence. I mean, it certainly is a consideration. I guess the thing to do would be...

Leo: But you've got to trust them.

Steve: As always, yeah, find some balance point, you know, for example, don't give no consideration to the security of the services that you're hiring. At least, you know, have them run the gauntlet and demonstrate that it makes sense for you to put some portion of your security in their hands because you are. You know, you are, when you're outsourcing a service, you're outsourcing the security of that service and that service's access back into your organization. And that's what bit the hospitals and bit all those dental practices when their common MSP got hacked.

So it's just - I sort of wanted to put it on people's radar to consider that, you know, if Dropbox hadn't been using CircleCI, well, they wouldn't have been prone to the CircleCI phishing emails. And so that couldn't have happened. Maybe something else would have happened. They would have gotten in some other way. But that's the way it happened. So it's very much like, you know, having exposed ports. Each of those things represent some exposure; and that means, you know, an expanded attack surface.

Two weeks ago - as we talked about last week when it was one week ago, now it's two weeks ago - the OpenSSL project maintainers told the entire world that one week from then a critical vulnerability would be patched and necessarily revealed to the world. So last week the severity, the good news was it was downgraded from critical to high. Since there is some possibility that one of the two problems could be weaponized, the advice remains that everyone using any v3.x.x of OpenSSL, where those x's aren't 0 and 7, which is to say if you're using anything before 3.0.7, which contains the two fixes, that should be looked at. So, okay. Here's what we know now, as I suspected last week, we would find out what was going on. Here's what the project maintainers wrote about the most serious of the two problems. It's got a CVE-2022-3602, now rated at high severity.

They said: "A buffer overrun" - which is of course where most of these problems begin. "A buffer overrun can be triggered in the X.509 certificate verification, specifically in name constraint checking." They said: "Note that this occurs after certificate chain signature verification and requires either a CA to have signed the malicious certificate, or for the application to continue certificate verification despite failure to construct a path to a trusted issuer." That meaning if it hadn't been signed. "An attacker can craft a malicious email address to overflow four attacker-controlled bytes on the stack. This buffer overflow could result in a crash, causing a denial of service" - meaning, you know, your service is denied because the thing crashed - "or potentially remote code execution. Many platforms implement stack overflow protections which would mitigate against the risk of remote code execution.

"The risk may be further mitigated based on stack layout for any given platform and compiler. Pre-announcements of the CVE described this issue as critical. Further analysis based on some of the mitigating factors described above have led this to be downgraded to high. Users are still encouraged to upgrade to a new version as soon as possible. In a TLS client, this can be triggered by connecting to a malicious server. In a TLS server, this can be triggered if the server requests client authentication, and a malicious client connects."

Okay. So the second of the two problems - there were two that were related. The second one is quite similar, but it only allows the attacker to overflow the stack with an arbitrary number of "dot," you know, period characters. I think that's hex 46. So the attacker's inability to overflow the stack with their own provided data, all they can do is dot characters, limits the practical danger to a denial of service that would result in a crash in OpenSSL. But the reason the more serious of the two was initially felt to be critical is that the stack overflow can be of attacker-provided bytes, for attacker-provided bytes. Which could be a jump or just enough code, for example, to elevate this task if it weren't already, or to bypass security checks, you know, whatever.

So what remains to be seen is whether anyone ever arranges to weaponize this attack. There's no doubt that many vulnerable instances of OpenSSL v3 previous to 07 will remain out in the world for the foreseeable future. They will have already been built into appliances that will never be updated. It's a relief that the trouble cannot be induced in an OpenSSL-based TLS server without the server first requesting a certificate from a client. That's unusual enough so as not to be a big issue.

But if an OpenSSL-based TLS client were to be induced into visiting a malicious server after this flaw were weaponized, that could result in the execution of code on the visiting client, thus compromising somebody who connects to a malicious server. And that could pose sufficient inducement to cause, that is, the potential of that could be sufficient inducement to cause major exploit creating players to investigate its weaponization. So we'll see if, a year or two from now, we're not talking about, whoops, remember that OpenSSL vulnerability that was downgraded to high, that should have been fixed wherever possible, well, you know, we'll see if that ends up happening. It could.

Okay. We're going to begin hearing of more instances of these sorts of reactions from the U.S. federal government; and, over time, it will become widely known that companies cannot simply ignore their security responsibilities with impunity. On Halloween, the FTC's Business Blog post was titled "Multiple data breaches suggest educational technology company Chegg [C-H-E-G-G] didn't do its homework, alleges the FTC." Now, we'll forgive the FTC for being cute about an educational company not doing its homework. But the points made in their blog posting about this were instructive.

The FTC wrote: "Chegg, Inc., sells educational products and services directly to high school and college students. That includes renting textbooks, guiding customers in their search for scholarships, and offering online tutoring. But according to the FTC, the ed tech company's lax security practices resulted in four separate data breaches in a span of just a few years, leading to the misappropriation of personal information about approximately 40 million consumers.

"The FTC complaint and some notable provisions in the proposed settlement suggest that it's time for a data security refresher course" - again with the educational approach - "at Chegg. Are there lessons your company can learn, the FTC posits or wonders, from where the FTC says Chegg failed to make the grade?"

Okay. Okay. In the course of its business - so here's what happened. California-based Chegg collected, they said, the FTC said, a treasure trove of personal information about many of its customers, including their religious affiliation, heritage, date of birth, sexual orientation, disabilities, and parents' income.

Leo: Why do they have my sexual orientation in the first place?

Steve: Exactly.

Leo: What the hell is that?

Steve: Exactly.

Leo: They're doing textbooks.

Steve: Yes. I know. Even the Chegg employee in charge of cybersecurity described the data gathered as part of its scholarship search service as "very sensitive."

Leo: Oh. Yeah. So you might - there might be a scholarship for queer scholars, something like that.

Steve: Okay.

Leo: So you'd have to give them that information, I guess, to find those scholarships.

Steve: In order to, yeah, to qualify, right.

Leo: It is. It's very sensitive.

Steve: Yes.

Leo: Yeah.

Steve: And four breaches. I mean, it's very sensitive, and they're not treating it responsibly. But wait till you hear, Leo, it's unbelievable. A key component of Chegg's information technology infrastructure was Simple Storage Service (S3).

Leo: Oh, boy.

Steve: Uh-huh.

Leo: S3 buckets can be secure, but they're often not. They're often not.

Steve: Cloud service offered by Amazon Web Services (AWS) that Chegg used to store a substantial amount of customer and employee data. The full complaint provides all the details, but the FTC cites a number of examples of what Chegg did and didn't do that were indicative of the company's lax security practices. For example, the FTC alleges that Chegg allowed employees and third-party contractors to access the S3 databases with a single access key that provided full administrative privileges over all information. Chegg did not require multifactor authentication for account access to the S3 databases. Rather than encrypting the data, Chegg stored users' and employees' personal information in plaintext.

Until at least April of 2018, Chegg "protected" - they have that in air quotes - passwords with outdated cryptographic hash functions. Until at least April 2020, Chegg failed to provide adequate data security training for employees and contractors. Chegg didn't have processes in place for inventorying and deleting customers' and employees' personal information once there was no longer a business need to maintain it. In other words, you know, it just kept accruing the data ad infinitum. Chegg failed to monitor its networks adequately for unauthorized attempts to sneak in and illegally transfer sensitive data out of its systems. In other words, across the board, your basic "do the minimum possible" laziness.

The report continues: "Should it come as a surprise that the complaint recounts four separate episodes that led to the illegal exposure of personal information? Incident 1 stemmed from a Chegg employee falling for a phishing attack that allowed a data thief access to the employee's direct deposit payroll information. Incident 2 involved a former contractor who used Chegg's AWS credential, the one credential, to grab sensitive material from one of the company's S3 databases, information that ultimately found its way onto a public website. Then came Incident 3, a phishing attack that took in a senior Chegg executive that allowed the intruder to bypass the company's multifactor email authentication system. Once in the executive's email box, the intruder had access to personal information about consumers, including financial and medical information. And Incident 4, a senior employee responsible for payroll fell for another phishing attack, thereby giving the intruder access to the company's payroll system. The intruder left with the W-2 information of approximately 700 current and former employees, including their birthdates and Social Security numbers."

Leo: Oh, god.

Steve: "In each of the four incidents cited in the complaint, the FTC alleges that Chegg had failed to take simple precautionary steps that would have likely helped prevent or detect the threat to consumer and employee data - for example, requiring employees to take data security training on the telltale signs of a phishing attempt." Because they fell for it four times, and nobody ever learned any lessons. No actions were taken as a consequence of those.

"To settle the case" - and, boy, have they gotten off easy - "Chegg has agreed to a comprehensive restructuring of its data protection practices. As part of the proposed order, Chegg must follow a schedule that sets out the personal information it collects, why it collects the information, and when it will delete the data. In addition, Chegg must give customers access to the information collected about them and honor requests to delete the data. Chegg also must provide customers and employees with two-factor authentication or other authentication method to help protect their accounts."

So it's going to get better, but this is just, you know, this is just a toothpick in a haystack; right? In this largely still unregulated industry, we're operating in a Wild West mode with nonexistent oversight until failures are egregious enough to bring governmental scrutiny. And how many of these incidents were caused by employees falling for phishing schemes? All four of them. Even an exec did. Yet there was no training provided. The reason is none of those breaches directly affected Chegg's bottom line. Oh, 40 million of their customers had highly sensitive data revealed? "Well, we're very sorry about that." Okay. Right.

Well, I'm not one who believes in government overreach and having Uncle Sam rummaging around in our private corporate businesses. But self-regulation isn't going to work here. One solution would be to only provide tools that provide security. Then at least security wouldn't need to be added on as an optional afterthought. But as we all well know, we're not there yet.

Leo: Everything you talk about on the show, Steve, is really a cautionary tale. And I just imagine these CISOs and CIOs and IT folks listening, going, oh, boy. Oh, boy. Did we secure our S3 buckets today? You know. This is good.

Steve: Well, and we talked a couple weeks ago there was some survey, it was IBM who did the survey, of the stress that CISOs...

Leo: Oh, can you imagine.

Steve: ...are under. I mean, it's just - it's not a - it's a horrible...

Leo: Tough job.

Steve: It's a tough - yes.

Leo: But a good job, important job. Thank you for doing it.

Steve: It needs to be done.

Leo: And we're glad you listen to Security Now! because that gives me some confidence that you're paying attention, which is good.

Steve: Okay. FinCEN, which is the U.S. Financial Crimes Enforcement Network unit which is part of the U.S. Treasury Department, published a 10-page report detailing ransomware-related events as reported by banks and other financial institutions through the Bank Secrecy Act (BSA). FinCEN said that in 2021, filings related to suspected ransomware payment substantially increased from 2020.

Okay. So we're nearly a year behind, right, because that's the way these reports go. Takes a while for them to filter through. So not like this year. We know this year was like a bang-up year, more so even than 2021. Anyway, 2021 substantially increased over 2020. 2021 saw a reported $1.2 billion in known ransomware payments paid out. The agency FinCEN estimates that roughly three quarters of these payments were made to ransomware gangs located in Russia. And of course that's all the ones that we're talking about, the big guys, all of this is Russian to a large degree. I've got a graph of the last few years of this. But basically it is your - it's not quite exponential, but it's more than linear. You know, it's more, yeah.

Leo: That looks like a hockey stick. It's a little hockey stick-y, yeah.

Steve: It's not good.

Leo: Going up fast.

Steve: So, boy. Yeah, we don't want Russia to be receiving our money. And the problem is when there's this much money behind it, $1.2 billion in cryptocurrency transfers, that's called incentive. And this is not what we want.

Leo: By the way, that's why it's so low in the left-hand side of the chart. You can really trace the success of ransomware to the rise of crypto.

Steve: Yes. Unless you could get paid without getting caught, there was really no way to make this happen. Remember it was Western Union transfers that was the way it was being done.

Leo: Yeah, or you'd go down and buy money cards from the 7-Eleven; right?

Steve: Right, right.

Leo: Sorry.

Steve: No, it's absolutely - it's been like the perfect storm where the bad guys realized, hey, this is great, we love this cryptocurrency stuff. Let's just ask for some bitcoin.

Akamai published their third quarter, their Q3 Threat Report for this year, 2022, which they released on right smack dab on the end, on Halloween. Since phishing has grown to become, by far, I mean, how many times have we spoken of it already in this 46 minutes, the most frequently detected first step in most successful attack scenarios. What Akamai's report had to say about phishing, I thought, was telling.

They said: "As covered in the Q2," that is, their previous quarter's 2022 report, "the overwhelming phishing landscape scale and magnitude is being enabled" - and this is news - "by the existence of phishing toolkits. Phishing toolkits support the deployment and maintenance of phishing websites, driving even nontechnical scammers to join the phishing adversary landscape and run and execute phishing scams." And anyone who's been listening to this podcast for long knows that's like the worst thing that we could hear, right, is you don't have to know anything now, increasingly, in order to pull off this, which is why there's so much of it.

They wrote: "According to Akamai research that tracked 299 different phishing toolkits being used in the wild to launch new attack campaigns, during the third quarter of 2022, 2.01% of the tracked kits were reused on at least 63 distinct days. 53.2, so a little over half of the kits were reused to launch a new attack campaign on at least five distinct days. And all 100% of the tracked kits were used on no fewer than three distinct days with the average toolkit reused on nine days during the third quarter of 2022." So the bad guys are being fickle about their toolkits. They jumping around trying different ones. And they're not - these are not long-lived campaigns. They're setting them up, sending out a bunch of emails, waiting for how long they would expect the email to take before somebody opened it and clicked on it. And they wait five, six, seven, eight, nine days, and then they go, okay, time to do a different campaign.

They wrote: "Further analysis on one of the most reused kits in the third quarter, counting the number of different domains used to deliver each kit, shows that kits that abuse Adobe and M&T Bank are top leading toolkits: Adobe with more than 500 domains" - just during Q3, I know - "and M&T Bank with more than 400 domains." Then they said: "The reusing behavior of phishing toolkits is more evidence of the trend of the phishing landscape that continues to scale, moving to a phishing-as-a-service model and utilizing free Internet services. Phishing attacks are more relevant than ever."

And it's interesting because their mention of utilizing free Internet services, remember, that was the one thing that the guy, the technical director of NCSC, who was the subject of last week's podcast, one of the things he said was I wish something could be done to limit free hosting services. That is where so much of the problem is. And at the same time he said, but what can you do in an open government...

Leo: Can't shut them in, yeah.

Steve: Exactly. But here, you know, utilizing free Internet services, the ability to just, you know, spin up free hosting and create free Internet service, that's a problem. So, but think about that, 299 distinctly different phishing toolkits. And as I said, what we've learned from observation is that the easier something is to do, the more it will be done. The Log4j vulnerability never swept the world as was originally feared because it turned out that the nature of the vulnerability meant that there was no one-size-fit-all exploit for it available. And if the script kiddies can't use something, then its use will be significantly curtailed. But if script kiddies can use something, then a feeding frenzy is the result. So on the front end it has never been easier to get into the phishing business. And on the back end, there's a huge market for the services of the so-called Initial Access Brokers; right? They're the ones who perform this, who develop initial access, and then resell it.

Leo: Right.

Steve: So any credentials that a phishing campaign can manage to obtain will find a ready market among those who can turn them into devastating network attacks.

I do have one little bit of news before I talk about Initial Access Brokers, and that is that Akamai reported seeing - although this was in their admittedly very skewed sample set, which I'll explain - they saw a 40% increase, from 25% to 65%, in the use of DNS over TLS. But that's not global. That's their enterprise and their own small and medium-sized business customers. But still, although this doesn't represent the world at large, currently more than 70% of all DNS remains over UDP. But what I think will happen is, this will be a very gradual change. As new systems are engineered from scratch, it's more likely that those new solutions will probably choose one of the encrypted forms of DNS, rather than old-school UDP. So we can hope. And it certainly says something that Akamai's own enterprise and small and medium-sized business customers really have started to adopt DNS over TLS.

Okay. As for Initial Access Brokers, another third-quarter report came out from a threat intelligence firm Kela, K-E-L-A. They published a report on the Initial Access Broker side of the network intrusion marketplace. Kela's report stated that during just this third quarter, this past third quarter that just ended this year, they found over 570 unique network access listings for sale, with a cumulative requested price of approximately $4 million U.S.

Okay. So just to be clear, someone responding and agreeing to purchase one of these 570 listings would be receiving, and this is something that's done through a Tor hidden service on the so-called dark web, they would be receiving the means to log into an unsuspected company's network with useful network privileges. Within that set of 570 listings, the average price to purchase access was $2,800, and the median price was $1,350. And prices have been rising since the second quarter. The total number of listings remained almost unchanged between the second quarter and the third quarter, appearing at the rate of around 190 new access listings per month.

So think about that. So there's a marketplace where people can go, and in fact as we'll get to it later, remember the numbskull Floridians, they actually went here, and they asked for access to CPA and tax preparer networks. I mean, this marketplace is that specific. You can go there, and you can say I want to get into the networks of these types of businesses, and you can purchase credentials that do that. And new credentials are appearing at the rate of 190 listings per month. That's 6.25 new listings per day, by the way. So anyway, and the average price, $2,800 to purchase access to somebody's network. And typically there's 570 of them up at any one time.

Wow. Okay. We will get to Florida in a minute. I found an interesting little bit that shared some details about how bank heists work. Although they don't receive a lot of coverage, over the past decade banks have not escaped ever-increasingly sophisticated cyberattacks. Many banks have been hacked and have collectively lost billions of U.S. dollars in serious intrusions. The two most notorious and successful threat actors that pulled off successful bank heists were a group called Carbanak, and also North Korea's Lazarus Group, which is an APT, an Advanced Persistent Threat group. Lazarus we've talked about before.

The attack geography, interestingly enough, has been evolving over time. Initial cyber heists tended to target organizations in North America and in Europe. Once those regions were fully explored, and security began tightening up, there was a move into Asia and Latin America. But as those banks also began to seriously upgrade their network defenses and security, movement has been now, more recently, in the direction of Africa, a region that has until now been left largely unscathed.

But a joint report published this week by security firm Group-IB and Orange's CERT team, a French-speaking cyber group tracked as, okay, we'll pronounce them "operator," although the "t" is a numeral "1," so "OPERA1ER," also known as Common Raven or the DESKTOP-group. They've recently been wreaking havoc across the African continent, well, recently from 2018 through 2021. This report covers nothing in this report since then. But actions have continued. The researchers said they linked this OPERA1ER group to 35 different intrusions at different organizations across 15 countries in Africa, with most of the attacks targeting banks.

Group-IB and the Orange researchers said that while the group used basic phishing attacks and off-the-shelf remote access trojans to gain an initial foothold in their victims' networks, once inside a network this OPERA1ER group has exhibited both restraint and patience. Some intrusions lasted for months, as the group moved laterally across banking systems, observing, mapping the internal network topology, and patiently waiting before springing their attack. The group's target was banking systems that handled money transfers. And this is what I found so interesting.

The report explained: "Once their network penetration had reached those most sensitive systems" - where the actual money transfers are managed - "the group would set a time for the heist and, working with a large network of some 400 money mules, would orchestrate a synchronized coordinated transfer of funds from the bank's larger legitimate accounts into the 400 mule accounts, with the money mules immediately withdrawing the stolen funds from their accounts via ATMs in a coordinated ATM cash-out before the bank's employees had the opportunity to react. The mules would refresh the ATM's screens at the appointed time, waiting for their account balance to suddenly jump up. Then they would drain the account for cash and quickly leave the area, thus of course bringing new meaning to the term 'decentralized finance.' The Group-IB researchers said they had linked OPERA1ER intrusions to bank heists totaling $11 million, but the group is suspected of stealing more than $30 million total, though not all the incidents have been formally confirmed."

So anyway, I thought that was interesting. The bad guys get in using phishing or remote access trojans, set up a presence in the networks, explore the networks, being quite patient, sometimes taking months until they determine what is there and get into a position where they're able to actually perform account funds transfers. They then reach out to their network, obviously a pre-established network of 400 individuals who then at a prescribed time go to ATMs where their own mule accounts have suddenly become wealthy, and dump all the cash out of the ATM that they can and then take off and head somewhere else. Wow.

Just to sort of keep an eye on DeFi, not to anyone's surprise, the DeFi platform Skyward Finance confirmed last Wednesday that a clever hacker had exploited a vulnerability in its smart contract system and made off with $3 million of cryptocurrency. And I guess at this point for us the proper expression would be, or the response, would be a yawn. And the DeFi platform Solend (S-O-L-E-N-D) said it lost 1.26 million worth of cryptocurrency following an Oracle attack on its platform which targeted the Hubble (USDH) currency. So it's hard to keep track of all these things these days.

Leo, you're going to love this one. In a big "What in the world took them so long?" bit of news, the Russian Ministry of Digital Development surveyed the country's largest IT firms, Russia's largest IT firms, to obtain their recommendations for the best replacement for Windows across Russian government and private-sector networks. The three contenders are all Linux-based operating systems, because what else could they be?

Leo: Yeah.

Steve: They are, I mean, you're right, there is nothing else.

Leo: Yeah, would they get Mac? No, of course not.

Steve: No, no. So they are the Astra Linux, ALT OS, and Red OS.

Leo: Red OS is the Chinese one; isn't it?

Steve: Oh, that's interesting.

Leo: China has its own Linux distribution the Chinese Communist Party recommends.

Steve: Yup, yup. It would certainly make sense that it was Red OS.

Leo: Red Linux, yeah.

Steve: And again, how many times have we, like, wondered, like what has taken them so long? Like how could Russia be using Windows? It's just astonishing to me.

Leo: They're often using pirated copies of Windows, and often using end-of-life pirated copies of Windows. So it's hideously insecure. The Chinese Linux is Kylin Linux, K-Y-L-I-N. And it's specifically for the mainland China market.

Steve: Well, and get this. It turns out that Russia would not have moved away from Windows but for their attack on Ukraine. Reportedly, the Russian government is seeking a replacement only now, after Microsoft pulled out of Russia, stopped delivering security updates to Russian systems, and started blocking Russians' access to Windows installation files. In other words, Microsoft left them with no choice.

Leo: Yeah.

Steve: And so, okay, Linux. Again, I wonder if - I guess I don't because they're moving to an open source operating system. Our NSA probably knows all about Linux, just as well as it does Windows. So it probably doesn't really make a difference one way or the other.

Okay, Leo, this one, wow. We've all seen war stories where, in the midst of battle, prominently marked Red Cross trucks come barreling in carrying noncombatants wearing wide Red Cross armband emblems with the hope and expectation that all combatants in the area, no matter whose side they're on, will respect the Red Cross's global neutrality and allow them to care for the wounded.

In a bizarre - and, okay, I was going to say interesting, but I think bizarre wins - move, they're trying to do this in cyberspace. After two years of study, last Thursday the International Committee for the Red Cross, the ICRC, has published their resulting report - again, took them two years - titled "Digitalizing the Red Cross, Red Crescent, and Red Crystal Emblems: Benefits, Risks, and Possible Solutions."

Okay. In explaining their intention, they wrote: "As societies digitalize, cyber operations are becoming a reality of armed conflict. A growing number of states are developing military cyber capabilities, and their use during armed conflicts is likely to increase. The ICRC (International Red Cross), has warned against the potential human cost of cyber operations and, in particular, the vulnerability of the medical sector and humanitarian organizations to cyber operations, both having been targeted in recent years.

"Against this background, the ICRC decided to investigate the idea of reflecting the internationally recognized distinctive Red Cross, Red Crescent, and Red Crystal emblems in the information and communication technology, i.e., a 'digital emblem.' Since 2020 the ICRC has partnered with research institutions to explore the technological feasibility of developing a digital emblem, and convened a global group of experts to assess its potential, benefits and risks. The idea and objective of a digital emblem was straightforward. For over 150 years, the distinctive emblems have been used to convey a simple message: In times of armed conflict, those who wear them, or facilities and objects marked with them, must be protected against harm."

Well, good luck. I wonder whether during these past two years of study those working on this have noticed how many hospital networks have been cyber attacked? You know, we're not dealing with declared hostilities in a battle theater where there's any sense of honor and conventions, Geneva or otherwise. I'll be interested to see how this one plays out. I mean, and what would prevent non-Red Cross organizations from putting up a Red Cross seal in order to protect themselves from attack? I mean, it's just loony. Okay.

Okay. Last Tuesday, the Department of Justice's U.S. Attorney's Office for the Middle District of Florida posted a press release with the title "Band Of Cybercriminals Responsible for Computer Intrusions Nationwide Indicted for RICO Conspiracy That Netted Millions." Okay. And, now, that's 36 millions, to be precise. Okay. The alleged tax fraud crimes took place between 2015 through 2019. DOJ officials said the group first purchased credentials from the dark web, allowing them to gain access to the internal networks of several Certified Public Accounting and tax preparation firms located across the U.S.

The group accessed the CPA and tax prep networks, stole the tax returns of thousands of taxpayers, created six tax preparation businesses in Florida and set up bank accounts and everything, I mean, full working businesses, and used those companies, those six tax preparation companies, to file more than 9,000 fraudulent tax returns in the victims' names and hijack tax refunds, directing them towards their own accounts.

And, surprise, surprise, somehow this was detected, and they didn't get away with it. Now they're all facing on the order of 20 years behind bars for RICO charges and fraud and money laundering and, you know, interstate felonies and you name it. I think what was most interesting and illuminating about this was the idea that I mentioned before that things are so well organized on the dark web that it's literally possible to search for network access by entity type. It's like, "Yeah, I'd like to purchase network access credentials for CPA and tax prep firms in the U.S. How much for how many?" Wow.

This piece from Microsoft, I'm not sure about this. Seems a little specious to me. It appears to be the month for reporting, and Microsoft is also out with their annual Digital Defense Report. The report contained a great many interesting tidbits, and buried among them was Microsoft's observation of an interesting change in China's profile. The observation begins with Microsoft noting that China's advanced persistent threat actors have leveraged significantly more zero-day vulnerabilities during the past year than anyone else.

Now, although most, if not all APT groups rely upon zero-day vulnerabilities for their exploits, Microsoft said that it had noted Chinese threat actors had an increased number of zero-days over the past year. And most interestingly, Microsoft believes that this sudden spike in zero-day exploits exclusively by Chinese threat actors is the direct result of a new law passed by the Chinese government last year. We talked about this last summer. The new law was passed in July of 2021, and it entered into effect in September of last year, 2021. It requires all Chinese security researchers to first report any new vulnerabilities they find to a state security agency.

And yes, at the time this did raise some eyebrows. It was roundly criticized within the security industry, while the Chinese government claimed that it only wanted to maintain an accurate catalog of vulnerabilities for the sake of making sure that local companies would not dodge responsibility for failing to patch vulnerabilities in time, thus leaving, obviously, Chinese users and government networks exposed to attacks. Uh-huh. Right. And that sort of sounds like a reverse-engineered rationale.

To put a point on it, the new law also contains several generically-worded clauses that could be interpreted to suggest that the Chinese government was setting up a secret process through which its offensive cyber units would have access to this trove of privately reported, at the time unknown vulnerabilities, while simultaneously suppressing the work of the infosec community for the benefit of the country's espionage operations. Although no solid evidence has come to light to support these theories, Microsoft appears to be sold on this narrative in its latest report.

They wrote: "This new regulation might" - this is Microsoft writing. "This new regulation might enable elements in the Chinese government to stockpile reported vulnerabilities toward weaponizing them. The increased use of zero days over the last year from China-based actors likely reflects the first full year of China's vulnerability disclosure requirements for the Chinese security community, and a major step in the use of zero-day exploits as a state priority."

To put a little more meat on the bone, Microsoft listed five specific zero-days as possible examples of abuse: two in Zoho ManageEngine, and one each in SolarWinds Serv-U, Atlassian Confluence, and Microsoft Exchange. Were exploits of these five zero-days developed by Chinese APT threat actors after they were reported through Chinese in-house vulnerability disclosure rules? We don't know. Maybe. On the other hand, would anyone be surprised to learn of zero-days in those applications? Hasn't all of that software been repeatedly plagued by major vulnerabilities and zero-day exploits discovered by other researchers and exploited by other threat actors? Of course. Of that there could be no doubt.

So perhaps a more accurate and rounded assessment would be that we cannot blame Chinese APT actors for looking at what everyone else is looking at and discovering the same zero-days that others are finding. Could they be getting a little help from the state's mandatory disclosure law? Again, maybe. But public evidence seems to be sorely lacking. What I wondered, like maybe reading between the lines, is whether Microsoft actually knows more than they're able to disclose without revealing their own sources and methods which they need to keep secret. Maybe this is a little bit of a shot across the bow saying read between the lines, China, because here's five zero-days that we think are suspicious. Maybe they have grounds, and they just can't talk about it.

Okay. So I love this idea. I'll be interested to see what feedback I get from our listeners because not everyone might like it. But it's interesting. The U.K.'s cyber group, the NCSC, will be scanning its public network space, looking for known vulnerabilities.

Leo: Hmm.

Steve: I think this is an interesting trend. We were of course just talking about the U.K.'s GCHQ NCSC cyber division last week when we covered the retirement of its technical director after his 20 years of service. And he certainly knew this was happening because this had to have been in the works for a while. So it was with interest that I noted what I think is the NCSC's excellent plan to periodically scan its own U.K. IP space searching for known vulnerabilities which are accessible on the public Internet and reporting them for remediation to the owners of those IP addresses. I think this is a terrific idea.

Okay. So they have an information page which they titled "NCSC Scanning information." It's not too long. I'm just going to share this because it's sort of in a Q&A fashion. They said: "This page provides information on the NCSC's scanning activities. You may have been referred here by information left by one of our scanning probes, if a system you own or administer has been scanned."

So they ask: "Why is the NCSC carrying out scanning activities?" They say: "As part of the NCSC's mission to make the U.K. the safest place to live and do business online, we are building a data-driven view of 'the vulnerability of the UK.' This directly supports the U.K. government cyber security strategy relating to understanding U.K. cyber risk. This will help us to" - three things - "better understand the vulnerability and security of the U.K., help system owners understand their security posture on a day-to-day basis, and respond to shocks, like a widely exploited zero-day vulnerability." That's interesting. So they'll be on top of that. When they find out something new like Heartbleed, for example, they would immediately scan the U.K.'s web servers and be proactive rather than passive.

Next question: "How does the NCSC determine which systems to scan?" They answer: "These activities cover any Internet-accessible system that is hosted within the U.K. and vulnerabilities that are common or particularly important due to their high impact. The NCSC uses the data we have collected to create an overview of the U.K.'s exposure to vulnerabilities following their disclosure, and track their remediation over time." Boy, this just sounds wonderful to me.

Next question: "How is scanning performed? To identify whether a vulnerability exists on a system, we first need to identify the existence of specific associated protocols or services. We do this by interacting with the system in much the same way a web browser or other network client typically would, and then analyzing the response that is received. For example, we may be able to determine the existence of a vulnerability known to exist in version X of a type of commonly used web server software by making a web request to the URL" - and then they give an example - ".../login.html and detecting the value 'version X' in the content of the page that is returned. If the vulnerability is then remediated in a subsequent version Y, we can identify this by similarly detecting the value 'version Y' in the response. By repeating these requests on a regular basis, we maintain an up-to-date picture of vulnerabilities across the whole of the U.K." Wow.

"What information does the NCSC collect and store? We collect and store any data that a service returns in response to a request. For web servers, this includes the full HTTP response, including headers, to a valid HTTP request. For other services, this includes data that is sent by the server immediately after a connection has been established, like the SMP headers, for example, or a valid protocol handshake has been completed. We also record other useful information for each request and response, such as the time and date of the request and the IP addresses of the source and destination endpoints.

"We design our requests to collect the smallest amount of technical information required to validate the presence/version and/or vulnerability of a piece of software. We also design requests to limit the amount of personal data within the response. In the unlikely event that we do discover information that is personal or otherwise sensitive, we take steps to remove the data and prevent it from being captured again in the future."

Question: "How can I attribute activity on my systems to NCSC Scanning?" They answer: "All activity is performed on a schedule using standard and freely available network tools running within a dedicated cloud-hosted environment. All connections are made using one of two IP addresses: 18.171.7.246 or 35.177.10.231." And they said: "Note that these IP addresses are also both assigned to 'scanner.scanning.service.ncsc.gov.uk' with both forward and reverse DNS records." So that's very cool. That means you could do a DNS lookup on scanner.scanning.service.ncsc.gov.uk, and it would return those two IPs. Or if you did a reverse lookup on either of those IPs, that's the DNS that you would get to know what that was.

They said: "Scan probes will also attempt to identify themselves as having originated from NCSC where possible, for example, by including the following header within all HTTP requests." And the header is X-NCSC-Scan: NCSC Scanning agent. And then they provide a URL to the page that I've been sharing so people can find out what that's about.

"What precautions and safety measures does the NCSC take when scanning?" They answer: "The NCSC is committed to conducting scanning activities in a safe and responsible manner. As such, all our probes are verified by a senior technical professional and tested in our own environment before use. We also limit how often we run scans to ensure we don't risk disrupting the normal operation of systems."

And finally: "Can I opt out of having servers that I own or maintain being scanned?" Answer: "Yes. Please contact scanning@ncsc.gov.uk with a list of IP addresses that you wish to exclude from any future scan activity, and we will endeavor to remove them as soon as possible once validated."

So, as I said, sign me up as a fan of this concept. Given the sad and sorry state of so much consumer crap and unfortunately the patch latency of so many enterprises, all of which is hung out on the Internet to be attacked, I think this makes a huge amount of sense. I mean, it's not like we're not all being scanned all over the place all the time anyway. I mean, I referred to it, it was one of the first acronyms or abbreviations that I coined, and that was IBR because I started getting involved in Internet security, and I thought, what is all this packet noise? And so it's Internet Background Radiation. It's just random crap out on the Internet that hits all of our IPs from time to time. So I think it would be great if the U.S. could take up similar responsibility and do something like this. Or maybe defer to individual ISPs to like police the traffic on their own networks and inform their customers.

Leo: Well, this was, you know, this was the big argument some years ago when spam - well, it's still a problem, but when it was really a problem. All an ISP would have to do is block port 25, the SMTP port, and they would effectively kill spammers on their network. And for a long time companies like Comcast, the biggest ISP in the U.S., wouldn't do it because they were afraid of the huge cost of tech support calls from people saying, well, I can't send my email anymore. And they eventually did do it. So ISPs, we've talked about this before. ISPs could, without doing the scanning that the British are doing, do a lot to police the outbound traffic from their networks.

Steve: Yup. And because they don't have to, they haven't done it.

Leo: Yeah. Yeah.

Steve: They have not been made to do it.

Leo: Yeah.

Steve: Yeah. And I think that their blocking of port 25 was also self-interest because they were getting complaints, like their network was sending all this spam.

Leo: Right.

Steve: And it was, yeah, it was a customer in their network. Cox, my cable provider, blocks port 25, so I have a way around that in order to contact my SMTP server at GRC. But something has to be done.

Just a quick note about Twitter since I'm about to share two listener feedback tweets. As my followers probably know, I have the blue verified check mark seal. And like so many others who have commented, I'm not going to be paying anything for it. I don't need any advanced features. I'm not paying anything for it now, and I'm certainly not going to be paying $100 per year to keep it.

Leo: Well, it would also devalue it because anybody who pays eight bucks regardless will get it. So it no longer verifies that you are who you say you are. It only means you paid eight bucks.

Steve: Right.

Leo: So it completely devalues - it doesn't mean verified anymore.

Steve: Yeah. So if it's taken away, I'll still be me.

Leo: Yeah, I'm not paying either. In fact, I got off Twitter. I'm done with that.

Steve: I did note one thing in passing which I thought was interesting. The Twitter alternative Mastodon reported that it had recently reached, not surprisingly, an all-time high of 655,000 active users after an influx of - get this - 230,000 new users just last week alone.

Leo: It's up to a million now.

Steve: Wow.

Leo: And I, you know, our server has a 7,000% increase in users, a 2,000% increase in interactions.

Steve: Wow.

Leo: You should join the - can I put a plug in for TWiT.social? I would love to have you. We even have, you know, on TWiT.social we have a custom icon that's your head. So I think you need to...

Steve: Well, all I really do with my Twitter account is tweet the link every week.

Leo: Yeah, and you don't have to give up Twitter to do that. But I suspect if you joined TWiT.social you would probably get in some very interesting conversations because people who listen to our show, many of them are there. And the thing to understand about Mastodon is, you know, I'm running - it's federated. So I'm running - it's like email. I'm running a server. But you can follow, and people can follow you from all over the Fediverse; right? If you were - and I will give this to you, @steve@twit.social, everybody would know to follow you. Or if you want to be SGgrc, whatever you want to be, you can be.

Steve: Well, I should be. I don't want to get engaged in conversation. That's not what I do.

Leo: You don't have to. It doesn't require it. It's up to you. I'm not going to push you into it, obviously. In fact, one of the great things about Mastodon, I'm a little reluctant to promote that we do this because I don't want a whole influx of Twitter people in here. I want people who, you know, are nice people.

Steve: Well, the good news is, Leo, the only people who are hearing this are the people who you do want.

Leo: Are nice people, yes.

Steve: Are nice people.

Leo: And that's a very good way of putting it, yes. It's a safe space here.

Steve: And that's how I feel about GRC's newsgroups. It's just it's a fabulous place where I'm able to get real work done. I should mention that I will be firing up a mailing list finally. I have to do it in order to announce SpinRite 6.1 to all of SpinRite 6.0's owners.

Leo: Ooh, exciting.

Steve: So that has to happen. So, and I'm going to - I'll create a number of different sublists and so forth. And I'm thinking as Twitter becomes sort of an uncertain deal, and frankly there are an awful lot of our listeners who are like, they've always refused to be on Twitter. So I will probably, one of the things that I'll do once I get a mailing system running is to just send out a short note every week, containing the show notes link because...

Leo: Oh, that's a great idea, yeah.

Steve: Yeah.

Leo: That's a great idea.

Steve: That way everyone will be able to get it. So, okay. Closing the Loop, two bits of feedback, as I said. I wanted to note that it was fun to receive all of the feedback from my discussion of my preferred keyboards last week, Leo. Not surprisingly, lots of people had opinions about keyboards. There's lots of discussions going on in various places now. So it turns out that I'm far from the only one who cares passionately about basically the way their primary device feels under their fingers.

David Stricker said: "This week you talked about ALT+TAB acting as MRU," right, Most Recently Used. He says: "But CTRL+TAB as round robin. Firefox has an option to set CTRL+TAB to act in MRU and is one of the main reasons I use it over Chromium-based browsers." He said: "I opened a bug with Chrome to allow MRU, and their response was simply 'Won't fix.'" So he said: "FF FTW." So anyway, I just wanted to share with our listeners something I never knew, which is that there was an option in Firefox that would allow you to change the behavior of CTRL+TAB so that it is not round robin, but MRU. And I would find that much preferable.

PCOwner said: "Steve, what is the best commercial cloud storage, secure, encrypted?" Okay, well, I know that there are many choices. But I did want to mention I am still, just to renew, still a fan of Sync.com, who I haven't talked about for a while. I've set up Sync to completely manage the file synchronization between my two locations, and it has never failed me. It's completely TNO (Trust No One) end-to-end encrypted. It has apps for iOS and Android, of course runs under Windows and Mac, presents a Sync directory under Windows and Mac, and allows for managed public link sharing despite the fact that it's end-to-end encrypted. So it has all the features that you would expect from a mature, secure, encrypted, commercial cloud storage provider.

What I did was to move a bunch of subdirectories that already existed on my system under Sync's automatically synchronizing Sync directory. So, for example, I have "c:\asm" where all of my assembly code work lives. So I moved that entire directory under the new Sync directory. Then I used Windows, there's a command in Windows, make link (mklink), which creates what's known as a junction point, you know, Linux refers to them as symbolic links or hard links. This creates a junction point where the relocated directory used to be at "c:\asm." This puts a link there so that all of the existing automation and batch files and everything that I have that expects my assembly language stuff to be at c:\asm, it's still there, as far as it's concerned, although it's actually under the Sync directory and now automatically synchronized between my multiple locations and available wherever I am.

The only feature missing, and they are painfully aware of it, is Linux client support. But I expect that their evaluation of the market for Linux, I understand it's a skewed demographic here in this podcast audience, but Windows and Mac have such a high percentage of the total desktop share that they don't seem to be making much headway on a Linux client.

Leo: No, because this has been going on for years.

Steve: Yes. And I did want to mention that without question for me, the best feature which I have used many times is that everything that is synchronized has full incremental versioning behind it, without the user ever needing to do anything. Boy, is that a win. And it has saved my bacon a couple times. I was once doing file versioning myself locally, but now it's just all built into the system that I'm using to synchronize my locations, and it's great. They have multiple plans, including a free 5GB plan that you can use to get your feet wet, and you can bump that, as I mentioned before when I talked about Sync, to a free 6GB if you go to Sync.com but use my affiliate code.

Actually, you can just go there in one jump. It's grc.sc/sync, grc.sc/sync. And that would give you an extra 1GB, and I get one added to my account, too. So anyway, still bullish about Sync. Again, I know that whenever I mention this, I get like 15 people all with different cloud sync providers, so I get it that there are alternatives. But this is the one that I can vouch for. And as I said, I've been using it, I use it every day, and it's never let me down.

And lastly - oh, boy, this is getting exciting - a quick update on where I am and what I'm doing when I'm not doing this podcast. I finished all of SpinRite's data recovery driver testing. All of it's working. The oldest drivers for BIOS-interfaced drives ended up needing a bunch of updating. That's all finished and tested. As the final piece of work, I turned my attention to SpinRite's command-line interface and its built-in command-line help. I updated everything in the online help with the new design. The redesign of the way it's going to work is finished, so the help guide is updated to reflect that.

Now I'm in the midst of rewriting much of SpinRite's command-line processor to make it, well, to bring it up to speed with all of the other changes that SpinRite has undergone. In the process of doing that, I needed to update SpinRite's "list" command which causes SpinRite to exit immediately after discovering and characterizing all of a system's mass storage devices which are accessible to it. It dumps that list in tabular ASCII text to the DOS console. For this new SpinRite, we also need a way of selecting drives through the command line. I could have just used the old way of indicating which line item in the listed table we wanted. But SpinRite power users use the command-line to automate SpinRite, and the ordering of drives could change over time if a drive was unplugged, or it went offline, or if a new drive was plugged into a lower numbered port, which would then get enumerated sooner and appear earlier in the table.

So a much more robust way of selecting drives is to allow a text match on any fields in the table. Since that includes the drive's model number and its serial number, it'll be possible to positively lock selections to specific drives. It'll also be possible to select multiple drives by class. For example, since one of the table's columns is "type," it'll be possible to give SpinRite the command "type AHCI," which will cause SpinRite to pre-select all of the system's AHCI drives, but none others.

So that's where I stopped working Sunday evening to put the podcast together. Tonight, well, probably not tonight because this is election night, so I will be in thrall. But tomorrow morning, first thing in the morning, I'll probably still have the election on in the background, but I'll be working on SpinRite, getting that finished and tested and then out into the hands of our group. So anyway, as I said to Lorrie during our walk yesterday, it's getting exciting. And we have - I think we're up to 406 registered testers in our GitLab instance. So we'll have a lot of people pounding on it, and we will move it as quickly as possible from Alpha into Beta. At which point I'll be able to make it available widely.

Leo: Yay. Is that it?

Steve: That's it.

Leo: There was, literally, something for everyone. I was waiting for you to talk about the guy who had a billion dollars in crypto in his coffee can in his backyard. Did you see that story?

Steve: I missed it.

Leo: He had stolen it. Let me see if I can find the details.

Steve: Oh, I did. I didn't know, I didn't realize it was stolen.

Leo: Oh, yeah, yeah.

Steve: I did hear something about someone who'd stolen a bunch of crypto.

Leo: Yeah. He'd stolen a bunch of crypto. And he put it on a little board because it's, you know, it strikes me you could just write down the number of your wallet. You don't need to actually...

Steve: Yes, you could, yes.

Leo: But for some reason he decided to put it on a board. Maybe he wasn't that sophisticated. Anyway, he had a billion dollars' worth of bitcoin. Was it a coffee can? Or it was hidden.

Steve: And he got found?

Leo: Oh, yeah. He got caught, and I think he's been arrested, yeah. Anyway, I don't have the - we'll probably talk about it on TWiT on Sunday because it's just a great story.

Steve: Yeah.

Leo: Mr. G. If you like what you hear here, you've got to check out his website, GRC.com. Yes, SpinRite's there, the world's finest mass storage maintenance and recovery utility. 6.0 is the current version; 6.1, as you heard, like just around the corner. You'll get it for free if you buy 6.0 now. You get an automatic upgrade. So it's worth doing that. You will want to have this software. If you have a hard drive or an SSD, you've got to have SpinRite.

While you're there, check out the show. Steve has two unique versions of the show, a 16Kb audio version and transcripts written by an actual human so they're actually legible. And you can use those to search or just read along as you're listening. He also has a 64Kb audio. GRC.com. You can leave him comments there. As you heard, he doesn't really want to talk to you. But if you want to leave a comment, go to GRC.com/comment, I'm sorry, feedback. Yeah. I don't blame you. I don't. I never read @ replies either. Or you can go to Twitter, @SGgrc. Sure I can't just sign you up, Steve, at TWiT.social? It'd be so much easier.

Steve: I do reply to DMs. I try to, you know, I mean, I'm present. But extended conversations, everyone would rather have SpinRite than me.

Leo: Yes, get to work. We have 64Kb audio. We have video, too, at our website, TWiT.tv/sn. There's a YouTube channel. You can subscribe in your favorite podcast client, as well, and get it automatically, the minute it's available. Some people like to watch live, like get the very freshest, hot off the podcast griddle version. We do the show Tuesdays. The time varies depending on how long MacBreak Weekly goes. Somewhere 1:30 to 2:00 p.m. Pacific's what we're shooting for, 5:00 p.m. Eastern, 22:00 UTC. Live.twit.tv is the stream. There's audio and video streams there. It's a nice thing to have in the background while you're working or whatever.

And if you're doing that, you might as well chat with us at irc.twit.tv. Club members can also chat in the Discord. And I guess, you know what, you could also comment on the TWiT.social there. Steve won't see it, but I will. Or on our Discourse, our forums at TWiT.community. So there's quite a few ways to interact, either synchronously or asynchronously, with me and other listeners. Don't expect Steve to get involved. He's got something better to do. More important.