

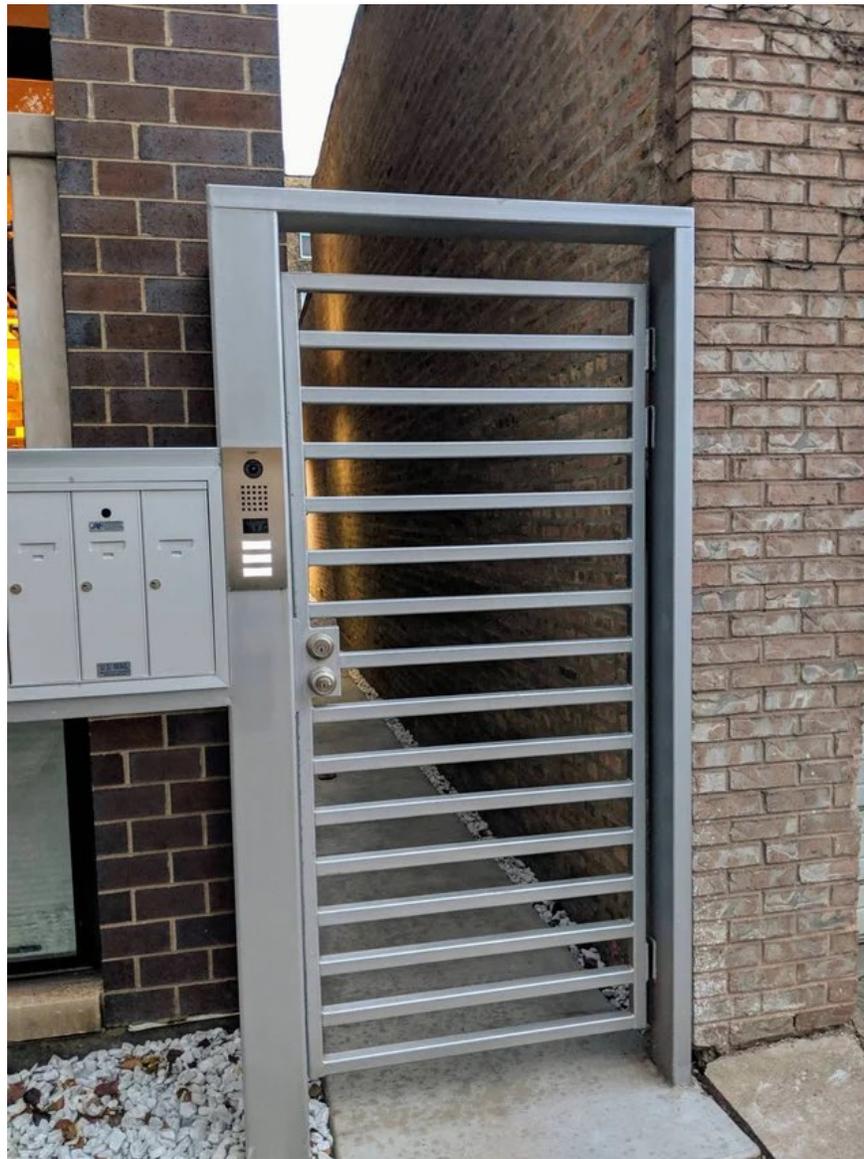
Security Now! #896 - 11-08-22

Something for Everyone

This week on Security Now!

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 0-days?, the NCSC will be scanning its citizenry... and more!

Can't get in? Hmm... How 'bout use the built-in ladder?



Security News

A minor DropBox breach

Last Tuesday, November 1st, Dropbox posted "*How we handled a recent phishing incident that targeted Dropbox.*" 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 a "not to worry" disclaimer:

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 have notified those affected and are sharing more here.

Skipping over a bunch of background, the part I wanted to share with our listeners was this:

At Dropbox, we use GitHub to host our public repositories as well as some of our private repositories. We also use CircleCI for select internal deployments. (CI stands for "Continuous Integration") 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 to login to CircleCI).

While our systems automatically quarantined some of these emails, 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 (OTP) to the malicious site. This eventually succeeded, giving the threat actor access to one of our GitHub organizations where they proceeded to copy 130 of our code repositories.

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.

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 be sure, we hired outside forensic experts to verify our findings, and reported this event to the appropriate regulators and law enforcement.

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 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. But those

filters failed just enough to allow bogus phishing attacks to reach their employees. And notice that these were code developers, not, for example, less sophisticated clerical or office workers.

The second point 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*" — the idea being that the more potential points of entry that exist, the greater the threat that any one of them 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 that the greater the number of ancillary services an organization employs, the greater is their phishing eMail attack surface.

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 sound: Why reinvent the same wheel over and over, especially where 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 the 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 that outsourced services model. And also remember all of those dental offices and hospital services that were hit with 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 that I wanted to make was where Dropbox wrote:

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 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 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 once 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, must be something that responsible organizations do. Dropbox's immediate response showed that they were ready and prepared for this eventuality.

OpenSSL follow-up

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. Last week the severity 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 of OpenSSL should arrange to update to 3.0.7 which contains the two fixes. Here's what the project maintainers wrote about the most serious of the two:

CVE-2022-3602 (OpenSSL advisory) [High severity] 01 November 2022

A buffer overrun can be triggered in X.509 certificate verification, specifically in name constraint checking. 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. 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) 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/compiler. Pre-announcements of CVE-2022-3602 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.

The second of the two problems is quite similar, but it only allows the attacker to overflow with an arbitrary number of "dot" (period) characters. So the attacker's inability to overflow the stack with their own provided data limits the practical danger to a denial of service due to a crash of OpenSSL.

What remains to be seen is where anyone ever arranges to weaponize this attack. There's no doubt that many vulnerable instances of OpenSSL v3 will remain out in the world for the foreseeable future. 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. And that could pose sufficient inducement to cause major players to investigate weaponization.

FTC sued and settled with a repeated offender.

We're going to begin hearing of more instances of these sorts of reactions from the US 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 posting was titled: "Multiple data breaches suggest educational technology company Chegg didn't do its homework, alleges FTC" We'll forgive the FTC for being cute about an educational technology 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 at Chegg. Are there lessons your company can learn from where the FTC says Chegg failed to make the grade?

In the course of its business, California-based Chegg collected 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. Even the Chegg employee in charge of cybersecurity described the data gathered as part of its scholarship search service as "very sensitive."

A key component of Chegg's information technology infrastructure was Simple Storage Service (S3), a cloud storage 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 of 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 multi-factor authentication for account access to the S3 databases.
- Rather than encrypting the data, Chegg stored users' and employees' personal information in plain text.
- Until at least April 2018, Chegg "protected" 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.
- Chegg failed to monitor its networks adequately for unauthorized attempts to sneak in and illegally transfer sensitive data out of its system.

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 Chegg employees falling for a phishing attack that allowed a data thief access to the employees’ direct deposit payroll information.

Incident #2 involved a former contractor who used Chegg’s AWS 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 and 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.

In 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.

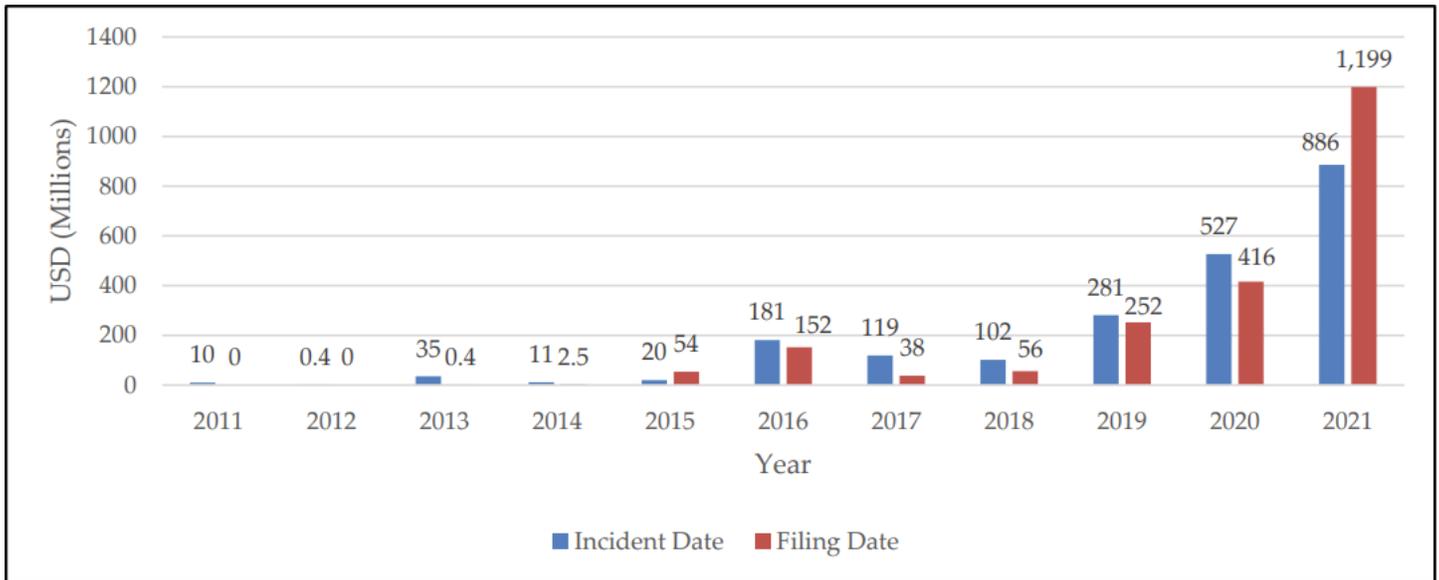
In each of the four incidents cited in the complaint, the FTC alleges that Chegg had failed to take simple precautionary steps that would have likely helped prevent or detect the threat to consumer and employee data – for example, requiring employees to take data security training on the telltale signs of a phishing attempt.

To settle the case, 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 that data. Chegg also must provide customers and employees with two-factor authentication or another authentication method to help protect their accounts.

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 those incidents were caused by employees falling for phishing scams? 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.”* I am not one who believes in government overreach and having Uncle Sam rummaging around in our private corporate business. 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 we’re not there yet, either.

\$1.2 billion in reported ransomware payment during 2021

FinCEN, the US Financial Crimes Enforcement Network unit which is part of the US Treasury Department, published a 10-page report detailing ransomware-related events as reported by banks and other financial institutions through Bank Secrecy Act (BSA). FinCEN said that in 2021, filings related to suspected ransomware payment “substantially increased from 2020,” amassing to \$1.2 billion. The agency estimates that roughly 75% of these payments were made to ransomware gangs located in Russia.



Akamai's Q3 Threat Report

Akamai's threat report for the 3rd quarter of 2022 was released on Halloween. Since phishing has grown to become, by far, the most frequently detected first step in most successful attack scenarios, what Akamai's report had to say about phishing was telling.

As covered in the Q2 2022 report, the overwhelming phishing landscape scale and magnitude is being enabled 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.

According to Akamai research that tracked 299 different phishing toolkits being used in the wild to launch new attack campaigns, in Q3 2022, 2.01% of the tracked kits were reused on at least 63 distinct days, 53.2% of the kits were reused to launch a new attack campaign on at least five days, and all 100% of the tracked kits were reused on no fewer than three distinct days with the average toolkit reused on 9 days during the 3rd quarter of 2022.

Further analysis on one of the most reused kits in Q3, 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 and M&T Bank with more than 400 domains.

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.

Think about that. 299 distinctly different phishing toolkits. What we have 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 feared, because the nature of the vulnerability meant that there was no one size fits all exploit available. And if the script kiddies can't use something, its use will be significantly curtailed. But if script kiddies CAN use something then a feeding frenzy will be 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 these IAB's — Initial Access Brokers. 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.

Encrypted DNS: In more welcome news, Akamai reported seeing a 40% increase, from 25% to 65% in the use of DNS over TLS within their admittedly skewed sample set, consisting of their enterprise and small and medium-sized business customers. So while this doesn't represent the world at large, where more than 70% of all DNS still remains over DNS, it's likely that as new systems are engineered, those new solutions will probably choose to use one of the several encrypted forms of DNS.

Initial Access Brokerages

While we're on the topic of Initial Access Brokers and 3rd quarter reports, the threat intelligence firm Kela also just published their report on the IAB side of the network intrusion marketplace.

Kela's report stated that during just the 3rd quarter of this year, they found over 570 unique network access listings for sale, with a cumulative requested price of approximately \$4 million US dollars. So, just to be clear, someone responding and agreeing to purchase one of these 570 listings would be receiving the means to log into an unsuspected company's network with useful network privilege. Within that set of 570 listings, the average price to purchase access was \$2800, and the median price was \$1350.

Also, prices have been rising since the second quarter. The total number of listings remained almost unchanged for the 3rd quarter, appearing at a rate of around 190 new access listings per month. But the total money requested during the 3rd quarter was considerably. This research didn't track the sales of this access, so we don't know what the turnover is, nor how long listings remain posted. But 190 new listings per month is 6.25 new listings per day.

How do today's bank heists work?

Though they do not receive a lot of coverage, over the past decade banks have not escaped ever-increasing sophistication of cyberattacks. Many banks have been hacked and have collectively lost billions of US dollars in serious intrusions. The two most notorious and successful

threat actors that pulled off successful bank heists were Carbanak and North Korea's Lazarus Group APT.

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

But according to a joint report published this week by security firm Group-IB and Orange's CERT team, a French-speaking cyber group tracked as OPERA1ER (also known as Common Raven or the DESKTOP-group) has recently been wreaking havoc across the African continent for the past four years from 2018 through 2021. The researchers said they linked the OPERA1ER group to 35 different intrusions at different organizations across 15 countries, with most of the attacks targeting African banks.

Group-IB and 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 victim's networks, once inside a network the OPERA1ER 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.

Once their network penetration had reached those most sensitive systems, 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 banks larger legitimate accounts into 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 jump up, then they would drain the account for cash, quickly leaving the area... thus bringing new meaning to the term "decentralized finance."

The Group-IB researchers said that they had linked OPERA1ER intrusions to bank heists totaling \$11 million, but the group is suspected of stealing more than \$30 million, though not all incidents have been formally confirmed.

De-Fi De-struction De-jour

Speaking of decentralized finance, 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 at this point the proper expression would be... <<Yawn>>

And the DeFi platform Solend said it lost \$1.26 million worth of cryptocurrency following an Oracle attack on its platform which targeted the Hubble (USDH) currency.

Russia moves to Linux

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 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? And they are “Astra Linux”, “ALT OS”, and “Red OS.” And as for what took them so long? They still wouldn't 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 other choice.

We're The Red Cross. Don't attack us, please!

We've all seen war movies where, in the midst of battle, prominently-marked Red Cross trucks come barreling in carrying non-combatants wearing wide red cross armband emblems with the hope and expectation that all combatants in the area will respect the Red Cross's global neutrality and allow them to care for the wounded.

In an interesting bid to move this idea into cyberspace, after 2 years of study, last Thursday, the International Committee of The Red Cross (the ICRC) has published their resulting report titled “Digitalizing the Red Cross, Red Crescent and Red Crystal Emblems: Benefits, Risks, and Possible Solutions.” In explaining their intention, they have written:

<https://www.icrc.org/en/document/icrc-digital-emblems-report>

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 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 (ICT), i.e. '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' are 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.

Hmmmmmm. I wonder whether during these past two years of study those working on this have noticed how many major hospital networks have been cyberattacked? We're not dealing with declared hostilities in a battle theater where there's any sense of honor and conventions, Geneva or otherwise. It'll be interesting to see how this one plays out.

Where there's a will there's a way

Get a load of this one: Last Tuesday, the DOJ'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." And that's \$36 millions to be precise.

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 (CPA) and tax preparation firms across the US. The group accessed the CPA and tax prep networks, stole the tax returns of thousands of taxpayers, created six tax preparation businesses in south Florida, and used those 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 they were detected and didn't get away with it. Now they are all facing on the order of 20 years behind bars.

I think what was most interesting and illuminating about this is the idea that things are so well organized on the Dark Web that it's possible to search for network access by entity type: *"Yeah... I'd like to purchase network access credentials for CPA and tax preparation firms in the U.S. How much for how many?"*

From China with Love

This 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 0-day vulnerabilities during the past year than anyone else.

Although most or not all APT groups rely upon 0-day vulnerabilities for their exploits, Microsoft said that it had noted Chinese threat actors using an increased number of 0-days over the past year. And most interestingly, Microsoft believes that this sudden spike in 0-day exploits by Chinese threat actors is the direct result of a new law passed by the Chinese government last year. We talk about this last summer. The new law passed in July of 2021 entered into effect in September 2021. It requires all Chinese security researchers to report any new vulnerabilities they find to a state security agency.

And, yes, this did raise some eyebrows at the time. 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 Chinese users and government networks exposed to attacks. Uh huh. Right. And doesn't that sound 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 unknown vulnerabilities while simultaneously suppressing the work of the infosec community for the benefit of the country's espionage operations. Though no solid evidence has come to light to support these theories, Microsoft appears to be sold on this narrative in its latest report.

Microsoft wrote: *"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 0-days as possible examples of abuse: two in Zoho ManageEngine, and one each in SolarWinds Serv-U, Atlassian Confluence and Microsoft Exchange. Were exploits for these five 0-days developed by Chinese APT threat actors after they were reported through China's in-house vulnerability disclosure rules? Maybe.

On the other hand, would anyone be surprised to learn of 0-days in those applications? Hasn't all of that software been repeatedly plagued by major vulnerabilities and 0-day exploits discovered by other researchers and exploited by other miscreants? Of that there can 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 0-days that others are finding. Might they be getting a little help from the state's mandatory disclosure law? Again, maybe. But public evidence is lacking. Perhaps Microsoft actually knows more than they're able to disclose without revealing sources and menthols that they need to keep secret.

The UK's NCSC scan plan

The UK's NCSC will be scanning its public network space for known vulnerabilities.

<https://www.ncsc.gov.uk/information/ncsc-scanning-information>

We were, of course, just talking about the UK's GCHQ NCSC cyber division last week when we covered the retirement of its technical director after 20 years of service. So it was with interest that I noted what I think is the NCSC's excellent plan to periodically scan its own UK 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 that this is a terrific idea and I'd love to see the US doing the same thing.

The NCSC's "NCSC Scanning information" page explains:

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.

Why is the NCSC carrying out scanning activities?

As part of the NCSC's mission to make the UK 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

UK Government Cyber Security Strategy relating to Understanding UK Cyber risk. This will help us to:

- *better understand the vulnerability and security of the UK*
- *help system owners understand their security posture on a day-to-day basis*
- *respond to shocks (like a widely exploited zero-day vulnerability)*

How does the NCSC determine which systems to scan?

These activities cover any internet-accessible system that is hosted within the UK 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 UK's exposure to vulnerabilities following their disclosure, and track their remediation over time.

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 ".../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 UK.

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 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.

How can I attribute activity on my systems to NCSC Scanning?

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
- 35.177.10.231

Note that these IP addresses are also both assigned to "scanner.scanning.service.ncsc.gov.uk" with both forward and reverse DNS records. 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:

X-NCSC-Scan: NCSC Scanning agent - <https://www.ncsc.gov.uk/scanning-information>

What precautions and safety measures does the NCSC take when scanning?

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.

Can I opt-out of having servers that I own or maintain being scanned?

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.

Sign me up as a fan of this concept. Given the sad and sorry state of so much consumer crap that's hung out on the Internet to be attacked, I think this makes a lot of sense and, as I said above, I think it would be terrific if the US could manage something similar.

Miscellany

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 paying anything for it now, I don't need any advanced features, and I will not be paying \$100 per year to keep it. If the blue verified check mark says, great. If it's taken away, I'll still be me. I did note one thing in passing which I thought was interesting: The Twitter alternative "Mastodon" reported that it had recently reached an all-time high of 655,000 active users after an influx of – get this – 230,000 new users last week alone. Yabba Dabba!

Closing The Loop

I wanted to note that it was fun to receive all of the feedback from my discussion of my preferred keyboards last week. It turns out that I'm far from alone in caring passionately about the feel of that primary interface device.

David Stricker @strickdd

This week you talked about Alt+Tab acting with MRU, but Ctrl+Tab is 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. I opened a bug with Chrome to allow MRU, and their response was simply "won't fix". FF FTW!

PCOwner @PCOwner2

Steve, what is the best commercial cloud storage, secure, encrypted?

I know that there are many choices. But I am still a fan of SYNC.COM, who I haven't talked about for a while. I've set SYNC up to completely manage the file synchronization between my two locations and it has never failed me. It's completely TNO end-to-end encrypted, it has apps for iOS and Android, presents a "Sync" directory under Windows and Mac, and allows for managed public link sharing. So has all of the features you would expect from a mature, secure, encrypted, commercial cloud storage provider.

What I did was move a bunch of subdirectories that already exist on my system under the auto synchronizing "Sync" directory. For example, I have "C:\asm" where all of my assembly code work lives. So I moved that entire directory under the new "Sync" pseudo-directory. Then I used Windows' make link (mklink) command to create a junction point where the relocated used to be, pointing to its new location. That way, all of my existing automation which expects c:\asm to be there continues to work without any change. And over time, as I have grown more and more comfortable with sync.com, I've migrated more of my volatile directories under it. I think that next time I'm setting up a system I'll put My Documents directory there.

The only feature missing, and they are painfully aware of it, is Linux client support. But I suspect that their evaluation of the market for Linux is so dwarfed by Windows and Mac that adding support is not high on their list. I do understand that might be a showstopping deal breaker for many users, but it isn't for me.

Without question, the best feature, which I have used many times, is that everything that's synchronized has full incremental versioning behind it without the user needing to do anything. Boy is that a win, and it has saved my bacon several times. I used to do file versioning myself locally. But now it's built into the system that I'm using to synchronize my locations.

They have multiple plans, including a free 5 gigabyte plan that you can use to get your feet wet. And you can bump that to 6 free gigabytes if you go to Sync.com by using my affiliate code which is (<https://grc.sc/sync>). You'll get an extra 1 gig, and I'll also get one added to my plan.

SpinRite

A quick update on what I'm doing when I'm not doing this podcast:

I finished all of SpinRite's data recovery driver testing. It's all working. The oldest drivers for BIOS-interfaced drives ended up needing a bunch of updating. That's all finished and tested. Then, as the final piece of work, I turned my attention to SpinRite's command-line interface and its built-in command-line help, updating everything there. The redesign is finished, so the help guide is updated to reflect that, and I'm in the midst of rewriting much of SpinRite's command-line processor. 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 form 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 went offline or if a new drive was plugged into a lower numbered port. 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 will 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 select all of a system's AHCI drives, but no others.

So, that's where I stopped working Sunday evening. This evening I'll resume that work to get that finished and tested. And then the work on SpinRite will be done, except for finding and fixing any bugs or things I haven't seen. Since this is a fully functional alpha release, I'll write some code for GRC's server, to allow for our licensed SpinRite testers to download their own DOS executable copies of the software, as is, and we'll start moving this from alpha to pre-release beta state.

And while our 401 GitLab-registered testers are pounding on SpinRite and logging any problems they find in GRC's GitLab instance, I'll switch my attention to the OnTime RTOS-32 to verify that it's the OS I want to commit to for the future. :)

