Security Now! #899 - 11-29-22

Freebie Bots & Evil Cameras

This week on Security Now!

What happens when you: • Run a Caller ID spoofing service? • Or when you mis-list and underprice online goods? • Or click on a phishing link for a cryptocurrency exchange? • Or consider working for a underworld hacking group? • Or use a webserver from the dark ages in your IoT device? • Or rattle your sabers while attempting to sell closed networking systems to your enemies? • Or decide whether or not to continue to suspend your Twitter ad buys? • Or login to Carnival Cruises with a passkey? • Or use hardware to sign your code? This week's podcast answers all of those questions and more!

Enough said.

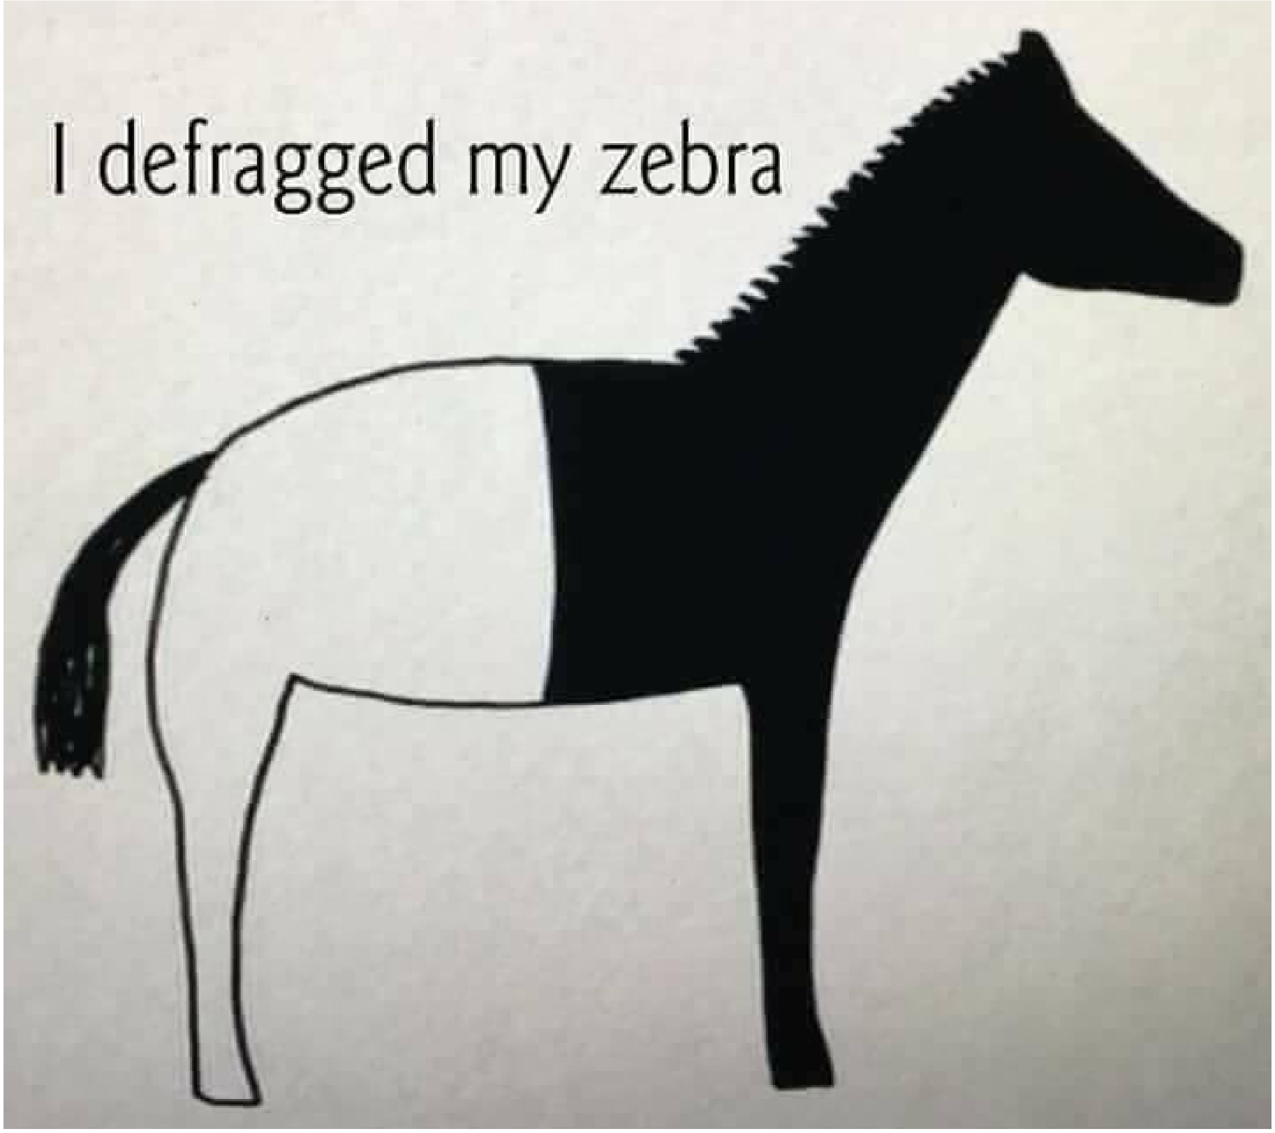

Security News

http://ispoof.cc/

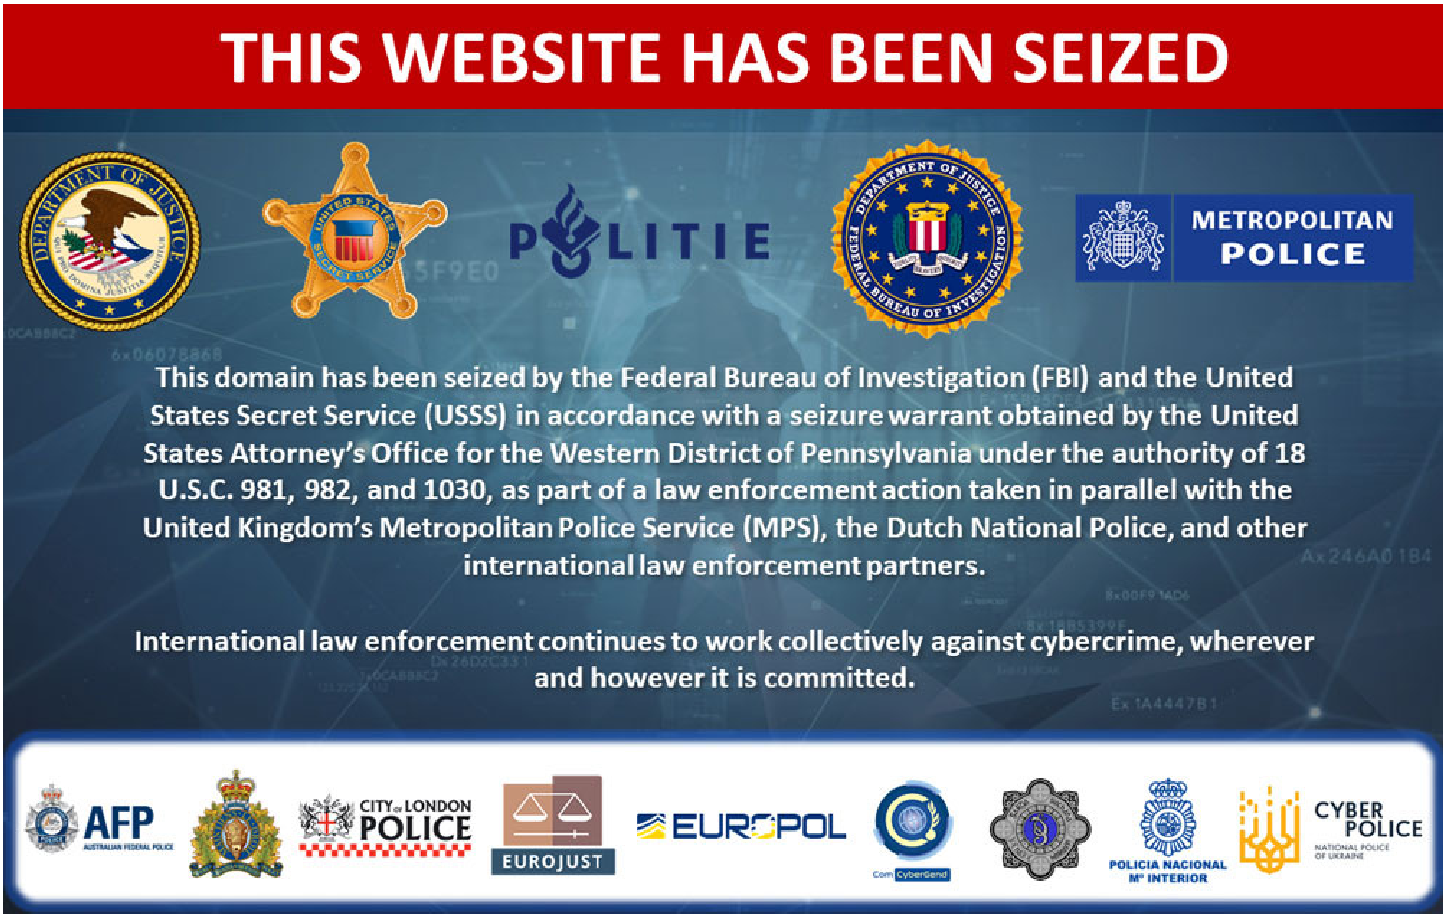

THIS WEBSITE HAS BEEN SEIZED

This domain has been seized by the Federal Bureau of Investigation (FBI) and the United States Secret Service (USSS) in accordance with a seizure warrant obtained by the United States Attorney's Office for the Western District of Pennsylvania under the authority of 18 U.S.C. 981, 982, and 1030, as part of a law enforcement action taken in parallel with the United Kingdom's Metropolitan Police Service (MPS), the Dutch National Police, and other international law enforcement partners.

International law enforcement continues to work collectively against cybercrime, wherever and however it is committed.

POLITIE
METROPOLITAN POLICE
AFP AUSTRALIAN FEDERAL POLICE
CITY of LONDON POLICE
EUROJUST
EUROPOL
CyberGend
POLICIA NACIONAL Mº INTERIOR
CYBER POLICE NATIONAL POLICE OF UKRAINE

iSpoof you no more

Get a load of this interesting bit of happening: Europol and law enforcement agencies from several countries including the FBI have seized the servers and websites of **iSpoof**, a service that allowed users to make calls and send SMS messages using spoofed identities.

The service launched in December 2020 and advertised itself as a way for users to protect their phone numbers and identities online, but Europol said that iSpoof was widely abused for fraud as it allowed cybercrime gangs to pose as banks and other financial organizations.

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

UK Metropolitan Police said that one hundred forty-two suspects were detained throughout the month of November, with more than 100 individuals detained in the UK alone, including iSpoof's administrators. Europol said iSpoof was being used to place more than one million spoofed calls each month, that administrators made more than €3.7 million, and that the service has been linked to fraud and losses of over €115 million worldwide.

UK police said they plan to notify all UK users who received spoofed calls made through iSpoof.

I was curious to see what the site looked like before the global takedown which displayed that site seizure page, above. So I turned to the Internet Archive project's Wayback machine and what I found just made my head shake:

https://web.archive.org/web/20220307052339/https://ispoof.cc/

The top of the site's very modern-looking home page proclaims: "Protect Your Privacy / Custom CallerID" – You can show any phone number you wish on call display, essentially faking your caller ID.

> *"Get the ability to change what someone sees on their caller ID display when they receive a phone call from you. They'll never know it was you! You can pick any number you want before you call. Your opposite will be thinking you're someone else. It's easy and works on every phone worldwide!"*

The most disturbing thing about this story is that the site was up and running for two years before it was brought down. That was a lot of damage done and you can imagine how word of mouth spread about this "handy" service among shadier types.

Here come the Freebie Bots!

What's a "Freebie Bot", you ask? A new class of Bot has been identified. And this one does something that would be difficult to predict, but once you hear what it does you think "Huh! Is that illegal?"

Last Tuesday, the anti-Bot research and security provider, Kasada, shared the results of its latest threat intelligence, which detailed the growing prevalence of so-called "Freebie Bots". Freebie Bots automatically scan and scrap retail websites searching for and purchasing mispriced goods and services, purchasing these discoveries at scale before the error is fixed.

Kasada research has found more than 250 retail companies recently being targeted by Freebie Bots, with over 7 million messages being sent monthly within freebie communities. (Just to be clear, these are not Furry communities, these are Freebie communities.) Members within one popular freebie community used Freebie Bots to purchase nearly 100,000 products in a single month with a combined retail value of $3.4 million. But Kasada's research revealed that due to significant mispricing, the total purchase cost of the goods for the Freebie Bot users was only $882. This allowed some individuals to realize a monthly profit of over $100,000.

The top items purchased using Freebie Bots during this time period included off-brand sleeveless halter neck mini dresses, Apple MacBook Air laptops, and deep cleansing facial masks. Many pricing errors were a result of decimal point misplacement, granting discounts as large as 99%. Using the speed and scale of a bot attack to rapidly purchase as much stock of these erroneously priced goods as possible, actors then turn around and resell the goods for a large profit.

So you can see how this could happen, right? Someone keying in a new item's retail listing gets into the habit of entering a decimal point before the last two digits of the price. But then they encounter a price formatted as a whole integer number of dollars without cents, and without thinking they place a decimal point before the last two digits, thus inadvertently reducing the listing’s price by a factor of 100.

It turns out that, at scale, across the entire Internet, these mistakes happen enough to have spawned the creation of a new class of Bots — automated retail mistake-finding bots — which will instantly purchase as much of something that's been mispriced as they're able to.

Human ingenuity knows no bounds. I suppose that while this might not be technically illegal, it certainly is unethical and dishonorable.

Anatomy of the real-time Cryptocurrency heist

The group PIXM Security, whose business is to protect end users from credential fraud, recently blogged about the details of an attack group they've been monitoring. The lengths this group will go to, to circumvent the newer “authorized device” protections which are becoming more common, is chilling, and I think it's worth having our listeners appreciate. The scammers will use an in-browser chat window to initiate a remote desktop session on the victims device, approve their own device as valid to access the users account, and then drain cryptocurrency from the victims wallet. Here are some of the chilling details...

When PIXMs Threat Research team first started tracking this group, the group was only targeting the Coinbase exchange. Over the last 30 days, the group has increased their coverage of cryptocurrency exchanges and wallets, to now include MetaMask, Crypto.com, KuCoin, and Coinbase. The spoofed domains are the typical slightly misspelled subdomains of azurewebsites.net.

The group employs effective second-factor relay interception when a user is spoofed into going to a look-alike site. Regardless of the credentials the user enters, whether they're legitimate or not, since the spoofing site cannot determine that, the user will be moved to a 2-Step Verification page after clicking ‘login’ where, depending on the platform in question, they will be prompted for either a 2-Factor code or their phone number used to retrieve their 2-Factor code. The criminal group will first attempt to relay these credentials and 2-Factor codes to the legitimate login portal associated with the platform they're spoofing. Once the user clicks ‘verify’ they will be presented with a message telling them unauthorized activity has occurred on their account.

As with the original Coinbase attack this group started with, this will initiate a chat window to keep the user on the phishing page in the event the 2-Factor code fails and the threat actor needs to start a remote desktop session with the victim to continue the attack. PIXM wrote that in their experience, regardless of whether the victim enters legitimate credentials or not, the group will ‘chat’ with the victim to keep them in contact should they need to resend the code or proceed to the second phase of the attack. The criminal gang's willingness to do this significantly increases the end-user's belief and engagement.

For a majority of the attacks this group carries out, they engage in direct interaction with the user. Their spoofed login and verification portals will, by default, return with a login error regardless of the actual standing of the user's account on the real exchange or wallet.

This process is intended to initiate a chat session with a member of the criminal group posing as a customer support representative from the exchange or wallet site being visited. The criminals will use this interface to attempt to access the users if their initial credential relay failed or time expired. They'll prompt the user for their username, password, and 2-Factor authentication code directly in the chat. The criminal will then take this directly to a browser on their machine and again try to access the users account. Should this also fail for any number of reasons (most common of which is that the device the attacker is using to access the victims account or wallet is not an 'authorized device' in the user's profile), the attacker will proceed to phase three with the victim. The group uses the "tawk.to" chat plugin on all of their sites and each with the same customer support representative named "Veronica."

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

The criminal has the user navigate to their email inbox associated with the crypto exchange or wallet account. They will instruct the user to login to their account on the exchange or wallet site. While the user is logging in, the attacker, who has control of the victim device, will enter a random character while the victim is entering their password, which will force it to fail. The attacker will click into the TeamViewer chat box without the victims knowledge, and ask them to enter their password again, which is just sending the password to the criminal in plain text.

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

The final draining of the user's cryptocurrency funds may be initiated during any of the previous attack phases since it is only contingent upon the attacker being able to successfully authenticate to the victims account from their own machine. Once the criminal is in the victims account, they will immediately begin transferring the cryptocurrency held in any of the victim's wallets to their own. They will keep the victim engaged and waiting as they steal their funds in the event that the service they are draining funds from might require some sort of email or phone confirmation of funds transfer. If that is the case, the attacker will assure the victim that this is normal and expected activity related to account restoration. Once all the funds have been sent from the victim to the criminals wallet, they will end communication with the victim having emptied their target's wallet.

Lookin' for something to do?
Karakurt, the group with known ties to former Conti gang members, and known for its hack-and-leak extortion operations, announced this week they are recruiting people to breach networks, malware coders, social engineers, and personnel to extort companies for payments.

Karakurt, which gets its name from a type of black widow spider, is not a ransomware gang. They don’t bother with encryption. They’re known for extortion and for demanding ransoms from $25,000 up to $13 million payable in Bitcoin. They don't target specific sectors or industries. The gang backs up their claims of stolen data with screenshots and copies of exfiltrated files as proof, and they threaten to sell or leak the data publicly if they don't receive a payment and Karakurt typically sets a one-week deadline to pay. Until they are paid they bully their victims by harassing their employees, business partners, and customers with emails and phone calls aiming to pressure the company into paying the ransom.

Their site on the dark web, a TOR hidden service, contains several terabytes of victim data, along with press releases naming organizations that had not paid and instructions for buying victims' data. resurfaced in May. The miscreants usually break into networks by either purchasing stolen login credentials; using third-party initial access brokers, which sell access to compromised systems; or by abusing security weaknesses in infrastructure.

Which brings us to their so-called "Great Recruitment" posting on the dark web. Since it was interesting and somewhat entertaining and I thought that it would bear sharing:

> *The Karakurt team is glad to announce some news. More than a year in private mode, but now we open the great recruitment! You can join our honorable mission — to make companies pay for the existing gaps in their cyber-security and for the inaction of the IT staff. So, our dear hack lovers, what we have for you:*
>
> - *Are you an experienced pentester and for some reason do not want to work with ransomware operators? You can find a better place in our team.*
> - *Do you work for a company that you hate with all your heart? Or maybe your boss fired you but forgot to turn off your network access? You can find solace in our arms.*
> - *You are a bearer of a sacred knowledge of malware coding? Disassembling? Exploit developing? The Karakurt team is ready to set interesting and non-trivial tasks for research, implementation of specialized software and modification of toolkits.*
> - *Are you from the financial industry? Do you know how to make money on quotes of companies whose shares are in poor condition? Know how to sell data in a specific market? We will hug you and love you more than anyone has ever loved you.*
> - *Are you from a data recovery company and know us? Let's be friends. Maybe even best friends.*
> - *Do you have social engineering experiences? There is also a vacancy.*
> - *Want to take revenge on capitalism through cyberspace? We will find you both a vacancy and a psychologist.*
> - *Perhaps you are a crazy researcher? We are really waiting for you, bro.*
>
> *The best hacker group, Karakurt, is waiting for you, our dear hack lover.*

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

Boa server vulnerability
The security firm Recorded Future found that a Chinese Advanced Persistent Threat Actor had leveraged a vulnerability in an IoT device to gain access to an electrical grid operator in India. And in a report last week, Microsoft said that they had identified the entry point for that attack. It was a tiny, somewhat obscure web server known as Boa which is said to be widely used across the IoT and ICS (Industrial Control Systems) space. http://www.boa.org/

As we all know, it can be very handy to have a nice simple and tight little web server. Such a tiny can be considered a component. Although Boa is written for Unix-like operating systems, it doesn't use the traditional UNIX fork and spawn approach of creating instances of itself to handle individual incoming connections. I didn't study Boa long enough to determine whether it's multi-threaded, spawning a new thread for each request. It might be purely serial. Since the UNIX / Berkeley sockets TCP/IP stack supports a queue of waiting connections, Boa might simply accept one connection after another using a single thread of execution. That would, indeed, make it quite lean. Apparently Boa is also quite fast.

All of that's okay. But here's the problem. It's not that Boa was first written and released 27 years ago in 1995. That's fine. The problem is that the last attention its source code received was 17 years ago, back in February of 2005.

In looking through Boa's development history, I noted with some interest that on October 4th and 5th of 2002 the Boa Developer's Conference was held. The official minutes of the event noted: “Larry and one of his sons stayed at Jon's house October 4-5, 2002. While the reasons were unrelated to Boa development, and in fact Larry and Jon spent only a few hours discussing Boa, computers, and the Free World, it seemed appropriate to refer to the event as a Developer's Conference. Here is a picture of Larry and Jon at Jon's house. (Left to right: Jon, Larry).”

Attendees Larry and Jon (at Jon’s house) during The Boa Web Server International Developer’s Connerence, 2002.

I have no doubt that these two have their hearts in the right place. But a web server they wrote 27 years ago and last tweaked 17 years ago, which has no support for secure connections, is currently in use, and apparently widely so, in, among other places, the operation of an electrical grid operator in India. Lord only knows where else this Boa Constrictor might be lurking.

Another of the multiple problems with today's software development is that it's too easy to grab this component or that component under the hopeful assumption – but with absolutely no proof – the result will be a secure system. We just assume things are secure until they're proven not to be, and then once again we hope for the best.

The dilemma of closed-source Chinese networking products

I dislike the idea of banning foreign companies from selling their products to whomever wants to purchase them. And the idea that networking and surveillance cameras of Chinese origin might incorporate designed-in Trojan capability seems far fetched to me. Presumably, such cameras are not phoning home to China but are networked locally. So the first instant that unexplained data was caught transiting the wire, there would be hell to pay. But at the same time, we can't prove the negative – we have no way of proving that there **isn't** any backdoor Trojan capability present in Chinese network and surveillance cameras. So I suppose that the recent actions from the US and the UK are understandable.

Last Friday the 25th of November, both the US and UK governments banned the use of Chinese networking and surveillance equipment, citing "national security"-related fears as the grounds for their decisions. The US Federal Trade Commission has banned the import and sale of networking and video surveillance equipment from Chinese companies Dahua (da-wow), Hikvision, Huawei, and ZTE. In the UK, the Parliament has instructed government departments to cease the deployment of security cameras from Chinese companies on "sensitive sites" such as government buildings and military bases. British officials said that Chinese-made security cameras should not be connected to core networks and that government departments should also consider removing and replacing existing equipment even before "scheduled upgrades."

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

Both Dahua (da-wow) and Hickvision had already lost a large chunk of their market in the US, after the US Treasury department sanctioned the companies for providing the Chinese government with facial recognition and video tagging solutions in the government's efforts to oppress the Uyghurs. And I recall that Hikvision was on our radar separately for something nefarious a year or so ago.

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

And the burden of trust is not symmetrical. Due to China's massive manufacturing and fabrication capability, they receive technology from the West, and the West purchases the resulting products from the East. Thus, more trust is required from the West than from the East.

So, I suppose my point is, we cannot discount such concerns as being purely hyperbolic and inflammatory. Our dependence upon our networks and digital infrastructure has been slowly but steadily growing for several decades. So it's only natural that at some point, someone at the national government level is going to wake up one morning and pose the big “but what if” question to their staff. It’s that “but what if” that was the driving factor behind the recent decision to “just say no.” Unfortunately, the resulting protectionism is both sane and rational.

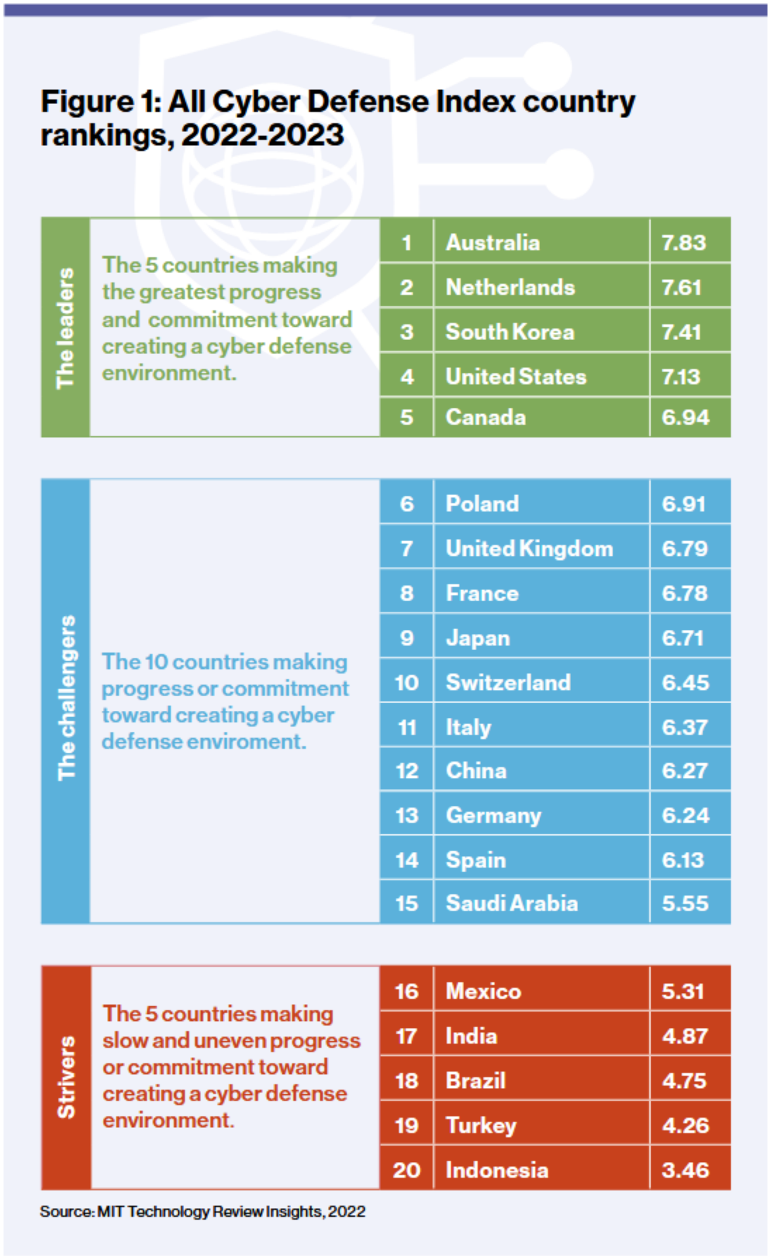

Figure 1: All Cyber Defense Index country rankings, 2022-2023

The leaders	The 5 countries making the greatest progress and commitment toward creating a cyber defense environment.	1	Australia	7.83
		2	Netherlands	7.61
		3	South Korea	7.41
		4	United States	7.13
		5	Canada	6.94
The challengers	The 10 countries making progress or commitment toward creating a cyber defense enviroment.	6	Poland	6.91
		7	United Kingdom	6.79
		8	France	6.78
		9	Japan	6.71
		10	Switzerland	6.45
		11	Italy	6.37
		12	China	6.27
		13	Germany	6.24
		14	Spain	6.13
		15	Saudi Arabia	5.55
Strivers	The 5 countries making slow and uneven progress or commitment toward creating a cyber defense environment.	16	Mexico	5.31
		17	India	4.87
		18	Brazil	4.75
		19	Turkey	4.26
		20	Indonesia	3.46

Source: MIT Technology Review Insights, 2022

The Cyber Defense Index

MIT recently published its rankings of national cyber defense. At the top of the list for the best defense is Australia. In second place it the Netherlands, 3rd place goes to South Korea, and the US beats Canada for 4th place with Canada in 5th.

Mit ranks those top five, followed by a middle group of ten, followed by the lowest five with Indonesia ranked dead last.

Malicious Docker Hub images

Just a quick warning to be careful of Docker images obtained from the official Docker Hub portal. The loud security firm Sysdig identified 1,652 malicious Docker images uploaded on the official Docker Hub portal. More than a third contained cryptomining code, while others contained hidden secret tokens that an attacker could later use as a backdoor into a server. Other Docker images contained proxy malware or dynamic DNS tools.

So, just be careful. Dockers are seductively easy to grab and deploy. But not everyone who’s creating and making them available is doing so out of the goodness of their heart.

Since we've been tracking 0-days for a while
I wanted to note that Google just fixed Chrome's 8th 0-day of the year. They updated Chrome to eliminate CVE-2022-4135 which was a heap buffer overflow (who's ever heard of such a thing?) which was found and exploited in Chrome's GPU component. The vulnerability was discovered by one of the Google TAG researchers and now it's history.

CISA on Mastodon
After a fake account was spotted for CISA's Director Jen Easterly on Mastodon, CISA now has an official account on the platform. The account is at the very popular infosec.exchange server where most of the industry's security researchers have been hanging their hat: https://infosec.exchange/@cisacyber.

Miscellany

This is not directly security or privacy related, but everyone's talking about Twitter and its uncertain future under the reign of Elon. I stumbled upon something that I thought our listeners might find as interesting as I did, because it appears to contain some actual facts. This is a note written by an unnamed executive director at an unnamed business-to-business organization. I presume it's anonymous because he would prefer not to have Elon Musk retaliate. The tittle of his posting was: "I told my team to pause our $750,000 per month Twitter ads budget last week" and he wrote:

> *I've seen a lot of technical and ideological takes on Elon Twitter* [I wonder whether that's a play on "Tim Apple"?] *but wanted to share the marketing perspective. For background I'm a director at a medium sized b2b tech company (not in financial services anymore) running a team that deploys about $80M in ad spend/year. Twitter was 8-10% of our media mix and we have run cost per engagement (i.e. download a white paper, register for an event, etc.) campaigns successfully since 2016.*
>
> *I had my team keep our [Twitter] campaigns live for 2 weeks post-takeover on the bet that efficiency would improve with fewer advertisers and that the risks were managed and probably overblown. I was wrong, and I think the things we saw in these last 2 weeks means many more advertisers will bail on the platform in the coming weeks (for non-ideological or virtue signaling reasons):*
>
> - *Performance fell significantly. CPMs didn't drop, but our engagement went way down. Maybe it's a shift in users on the platform, maybe it's ad serving related.*
>
> - *Serious brand safety issues. Our organic social and CS teams got dozens of screenshots of our ads next to awful content. Replies to our posts with hardcore antisemitism and adult spam remained up for days even when flagged.*
>
> - *Our entire account team [at Twitter] turned over multiple times in 2 weeks. We had multiple people (AE, AM, analyst, creative specialist) supporting our account and they all vanished without so much as an email. We finally got an email with a name for an AM last*

> *week but they quit and we don’t have a new one yet.*
>
> - *Ads UI is very buggy and login with SSO and 2FA broken. One of my campaign managers logged in last week and found all our paused creatives from the past 6 years had been reactivated. Campaign changes don’t save. These things cost us real money.*

Since I hadn’t encountered anything as substantive as that, I thought that it was interesting to see and understand a bit about what’s going on from the perspective of one of Twitter’s advertisers who views the service dispassionately and doesn’t care one way or another who’s doing what. Twitter either will be or won’t be a means to their ends.

And in a related piece in a security newsletter I recently scanned, the statement was made

> *“Some threat intelligence companies are telling their customers that they can no longer guarantee takedowns of malicious or reputation-damaging content from Twitter as there is nobody in Twitter's abuse team to respond to requests anymore.”*

Closing The Loop

KerryOnAnon @MrIndigo_

> *Hi Steve! Finally listening to the latest episode #898 and I started wondering, is quantum computing going to be just a faster way to guess passwords or is there another attack vector? In other words, is it just going to be a faster way to brute force attack passwords?*

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

The important thing to understand here is that **some** of today’s crypto, but only some of it, depends upon the traditional, time-proven difficulty of factoring a very large number into its two, half as large prime number components. That's it. That's all that the fervor surrounding quantum computing is about. The ability to do a couple of things quickly that are currently insurmountable. But it's only the asymmetric key crypto that quantum computing might be able to someday weaken. NONE of the other crypto that we also depend upon today will be affected. Symmetric key crypto, like our beloved AES ciphers or today's strong hashing algorithms will not be affected at all. And won't need to be changed.
I was thinking about quantum computing, and looking for a good analogy of the effort, the promise and the difficulty that it presents. And what popped into my head as being in almost every way similar, was power generation, at scale, via nuclear fusion. It's a useful analogy. It requires crazy, way out there, new physics and new materials and new technologies. And, like quantum computing, we've been chasing fusion for decades, driven by the promise of “what if.” Just like quantum computing. And incredible amounts of ingenuity and money have been sunk into it. Many different approaches have been tried and discarded. And yes, we're creeping

forward little by little, inch by inch, tantalizingly. Just enough to keep the investment cash flowing. But boy, is it difficult! In order to fuse matter, we must create, contain and compress the hottest plasmas humans have ever handled, and at this point it's as much art as science.

Will we get there someday? Maybe. Maybe not. It's still not clear. But as with quantum computing, we do appear to be making some progress year after year, learning as we go.

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

A listener who requested anonymity

> *Hi Steve. In the last episode of Security Now, you talked about passkeys.directory, which lists web applications that support Passkeys. I wanted to share my observations with you.*
>
> - *The website owner chose to manage it with no transparency. When I saw it, I thought there must be a Git repo where I could open an issue or a change request. Surprisingly they choose to use Google Forms, which masks all the review and approval processes.*
>
> - *I've noticed that many companies in this list are also customers of OwnID, which is listed as an Authentication Provider, including Carnival Cruises.*
>
> - *Investigating the OwnID flow:*
>
> *When Leo pressed the fingerprint button, the QR code encoded a URL that sent his iPhone to passwordless.carnival.com with a session identifier. Then he performed a WebAuthn authentication on his iPhone. Once completed, the session got updated on the server and the browser on his laptop logged in.*
>
> *This flow is using WebAuthn's Passkeys, but not likely the way it was designed to be used:*
>
> *WebAuthn Phishing Resistance mechanism works in a way that a JaveScript API called on the browser triggers the underlying CTAP2 library and matches the domain a key was registered in and the domain asking to authenticate. By implementing WebAuthen as it is in Carnival, the Phishing Resistance mechanism suffers from a flaw.*
> *As an attacker, you can spoof Carnival's login page, so the user sees the same page, only a different domain. When you click the Biometrics button, the attacker's backend will send a request to Carnival to get a QR which encodes the passwordless.carnival.com. Then the phone would ask you for your face or fingerprint to authenticate with a Passkey, which will update the session on the backend, and the attacker gets in.*
>
> *The right way to implement Passkeys is by calling the WebAuthn API on the laptop's browser*

(instead of presenting the QR that will open a browser on the mobile phone) and letting the browser do its job, presenting native WebAuthn screens, including a QR which is scannable from a mobile device. This way, the domain you're authenticating to is passed in a side channel (Push/BLE) to the mobile phone directly from the browser, and a phishing site will be blocked as the credential on the phone was registered under the original domain.

This listener is 100% correct. And, by the way, he's a developer for an authentication provider.

Another way to say this is that rather than doing the work of upgrading their own servers to become a first-party passkeys provider, Carnival Cruises has outsourced their authentication responsibility to a 3rd party provider, in this case OwnID. But in doing so, by punting in this way, they've bypassed passkeys phishing protections. This gives their visitors the false belief that they're getting the hack-proof benefits of passkeys without actually having them. This could be transient. But OwnID will presumably be selling their "instant onboarding" services and most websites will simply want easy logon without really caring about their visitor's security.

Christopher Ursich / @chrisursich

SN topic request: Hardware Security Modules (HSMs)
You said you had one. Besides the technical crypto, can you describe how you interact with it in practice to sign your code?

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

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

Part of the effort toward the end of the work to publish the final SpinRite v6.1, which will be a hybrid DOS and Windows app, will be automating this code signing process, since each owner's copy embeds their license information, making their executable unique. So each one needs to be individually signed on-the-fly by the server, as it's downloaded.

What's going to be really annoying is that Windows Defender will always be complaining, for every single user, that the user-specific custom SpinRite file is not commonly downloaded, thus needlessly warning and alarming its users. We've seen that no degree of reputable signing is able to bypass that alarm.

Dangard @Dangard

> *Steve, How can I get access to test the pre-release version of SpinRite 6.1? Feel free to email me or just respond here. Thanks so much for your work on SpinRite. I have drives waiting for 6.1 :)*

sdholden / @sdholden

> *Hey Steve. Not sure the best way to reach you about the git server for SpinRite, so thought I 'd start here. When I try to create an account, I get a dialog box asking me to sign in instead of allowing me to create a registration...?*

To both listeners and everyone else:

In case some of you hadn't noticed, the Internet has, sadly, become a sewer full of both Bots trolling constantly and even human labor farms being paid for creating accounts online. I've been running two web forum servers for years. Despite having all manner of entrance barriers erected, like requiring the correct answer to the question "What software is Steve best known for?" in order to create an account, five out of six of the account registrations were bogus in those forums. At one point we had more than 6,500 users registered in GRC's forums. Now it's a bit more than 1100 after I spent several days working to get that under control. I erected much tougher barriers and I have mostly gotten in under control, and since I erected those stronger barriers, 20,204 additional account creation attempts have been thwarted. The reality is that, today, running any sort of open web service results in a torrent of bogus registrations. And even with all of that in place, the wonderful volunteer moderators I have, who make time to read everything, are still removing users who attempt to subtly pollute our content.

GRC's forums need to be open, so I have no choice other than to erect the strongest account creation barriers that I can, then apologize to those whom we mistakenly reject as false positives and also weed out those who slip past our barriers due to false negatives.

But GRC's GitLab server does not need to be open to all. So it's closed. Its account creation page is protected by a magic incantation which must be provided before the troll that guards the bridge will allow newcomers to pass. It requires insider information which can only be obtained by participating in GRC's old school, blessedly wonderful text-only NNTP newsgroups. Once someone shows up there, and is able to post, they can ask how to satisfy our cantankerous GitLab troll. But also note that we're not using GitLab for any social interaction. We're only using it for issue management.

At this point, what I need is feedback from people who are testing SpinRite v6.1. Since we have a handful of known issues to fix (I'll get to that in a moment) it's best for newcomers to join and catch up on all of the various threads, in order to eliminate duplicate postings of already-known issues.

So if anyone is really and truly interested in participating in SpinRite v6.1's testing, you are invited to head over to GRC's "discussions" page and create a connection to our news server. Find the grc.spinrite.dev group and say hi!   And speaking of SpinRite... It's working!

SpinRite

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

A few hours after my first release announcement, **DarkWinX post on Friday at 2:44pm:**

> *Well I can already report success with a USB. In my race to find something to eagerly test on, with the short time I had, I grabbed an old USB I received with the purchase of Starcraft 2. I figured I'd reformat it with InitDisk and run SpinRite from there. So I put it in the computer and started Initdisk. It waited, and waited, for about 30 seconds. Eventually the USB was recognized and showed up in Windows and I could nuke it. I tried it again and it still took around 30 seconds to load.*
>
> *So I figured, maybe not the best USB to run SR from. So I found another. I thought why not run SpinRite on the problem USB as a target - so that's what I did. After a Level 2 scan, without finding anything wrong, I rebooted, plugged in and.... instant success. That USB now loads inside Windows instantly every time. Looking forward to testing some more!*

Saturday morning @ 8:39, Mark Ping posted:

> *Finished the level 2 in 2 hours for 1 TB. Then ran level 4 and it took 9 hours, 37 minutes for 1 TB, compared with 150 hours before. Spinrite is back baby!*

Dale F., Saturday evening at 10:12 posted:

> *I have a 500MB laptop drive that I put in a SABRENT portable enclosure. After I dropped it about 2 years ago, it could not be recognized by any PC — or by SpinRite 6.0. So I said to myself “just have to wait for 6.1.” On Friday, I ran a level 2 with SpinRite's first alpha release, and 1 hour later, it was good as new. THANKS, Steve!!*

Frankly, SpinRite's first functional pre-release debut could not have gone much better, and it went far better than it might have. Over the weekend, using the feedback provided by the large group of avid testers, we moved SpinRite through three more releases to its 4th alpha release by mid Sunday afternoon and, with only a few exceptions, it's working well for everyone.

Overall, it's 100% functional in every way that matters. There are a number of things that I need to fix, like SpinRite's various clocks are not continuing to operate while it's deep into data recovery. I recently rewrote that entire data recovery system, and I just forgot to periodically update the clocks. And I'm going to change the entire way that works so that it's much better.

Another example is that SpinRite's predictions of its remaining time to run is not working right when it is started midway into a drive rather than somewhere later in the drive. It was working once and something I did broke it.

So, right now the newsgroup gang is continuing to pound away on alpha 4, logging everything they encounter in our GitLab instance. While that's underway my own highest priority is to make a decision about that next operating system that I'm considering purchasing and moving to. Its licensing deadline is the end of the year. I expect that to take a few days, then I'll return and work to get SpinRite's DOS executable completely finished.

One thing happened this morning that completely caught me off guard. I hired Greg 32 years ago tomorrow. So tomorrow he will have been providing technical support for SpinRite for 32 years. Yesterday he fired up the latest SpinRite v6.1 alpha and ran it on a bunch of drives he had around. He said that he ran in on a 1TB spinner which took about 2 hours. That beats two weeks, and it still wasn't instantaneous. But then he scanned a 128 gigabyte SSD in 5 minutes, and he was stunned. So he told me on the phone this morning that the number one question he was certain people were going to be asking, once SpinRite's previous users started using 6.1, was how could 6.1 possibly be so much faster?