Security Now! #917 - 04-04-23

Zombie Software

This week on Security Now!

This week we answer questions which arose during the past week: When is an attack not an attack? When our AI overloard arrives how shall we call him? Why has Italy said NO to ChatGPT? What does Twitter's posting of its code to GitHub tell us? Why is India searching for commercial spyware less well know than Pegasys and what does the Summit for Democracy have to say about that? Has the FDA finally moved on the issue of medical device security updates? And seven years after the first "Hack the Pentagon" trial, the Pentagon remains standing, or does it? Then, after addressing a quick bit of miscellany, listener feedback and an update on my ongoing work on SpinRite, we use CISA's KEV database to explore the question of how exactly we define "Zombie Software" and answer the question of whose brains will the zombies eat?

The Dangers of Legislating Technology

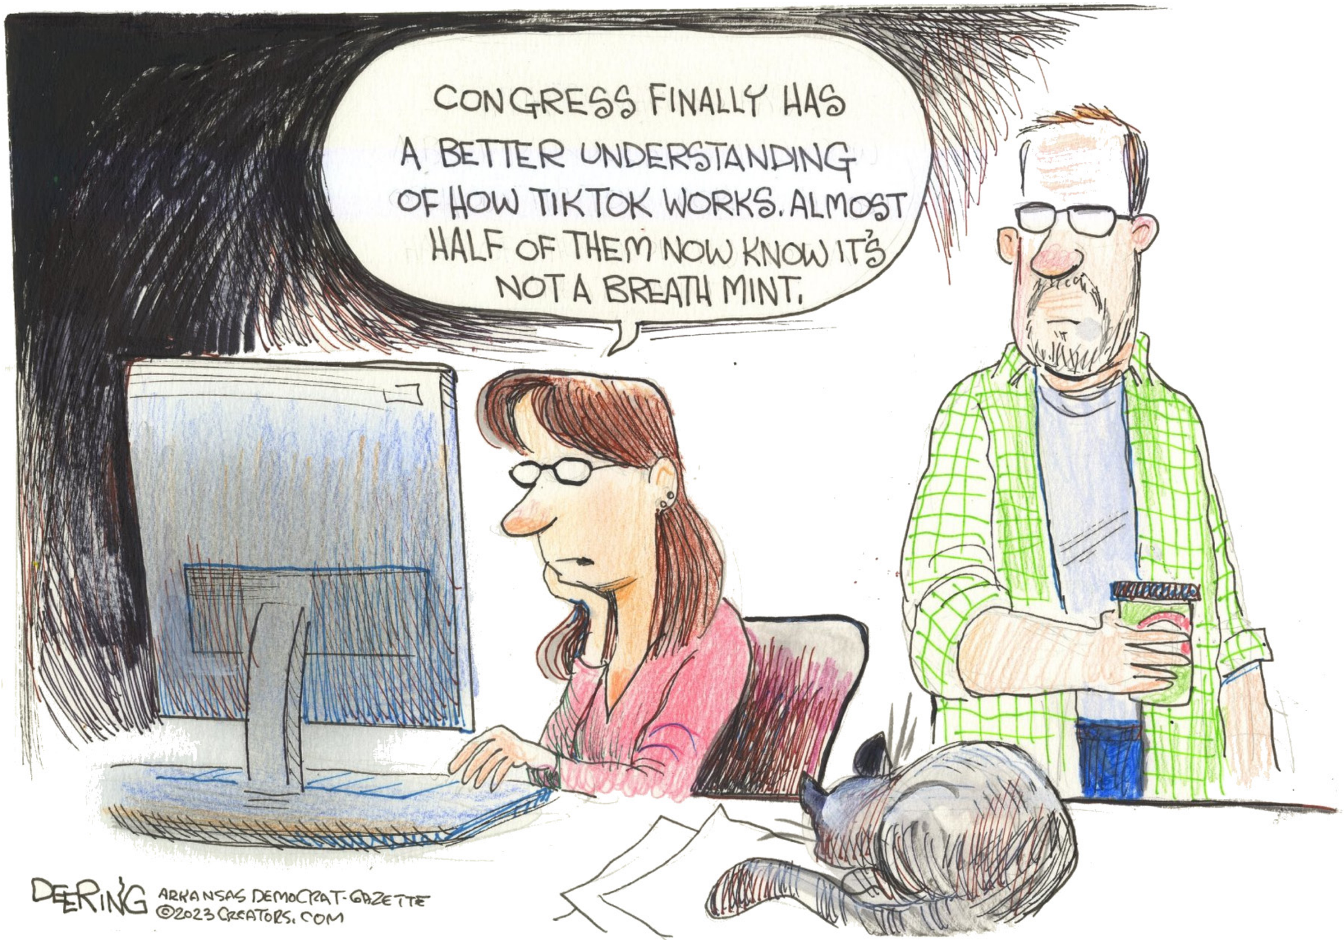
CONGRESS FINALLY HAS A BETTER UNDERSTANDING OF HOW TIKTOK WORKS. ALMOST HALF OF THEM NOW KNOW IT'S NOT A BREATH MINT.
DEERING
ARKANSAS DEMOCRAT-GAZETTE
©2023 CREATORS.COM

Security News

So... Not an attack, then?

Recall two years ago all of the fuss, bother and more than a little bit of horror over the attack at a Florida water treatment plant in Oldsmar, Florida? The report that the world received was that a caustic and potentially poisonous concentration of lye had been dumped into the facility's municipal water treatment tanks — remotely over the Internet — by a disgruntled former employee whose remote access credentials had, we were all told, not been canceled and deleted after this departure.

And the story had so much flair and detail. Recall that the reports of the incident described how a worker at the plant saw his computer being remotely accessed and controlled. His mouse moved to open functions to control the plant's water treatment protocols, and then the amount of sodium hydroxide, or lye, in the water was changed from about 100 parts per million to 11,100 parts per million. Whereupon the operator immediately reduced lye's level to the proper level and alerted his supervisor.

And everyone got in on the act. The hack gained worldwide notoriety after the local Pinellas County Sheriff held a press conference which, in turn, prompted an investigation led by the FBI and the U.S. Secret Service, as well as a joint federal advisory warning water treatment facility operators of the dangers they faced from hackers and urging them to upgrade their security systems.

Gristy as that story was for our mill, what if that's not what happened at all? That's right. According to former Oldsmar City Manager Al Braithwaite, who was with the city at the time, the incident was not a hack at all. It was just a case of an employee mistakenly clicking the wrong button, then alerting his superiors to his error. Former City Manager Braithwaite now describes the incident as a total "non-event" which was resolved in two minutes, but said law enforcement and the media seized on the idea of a cyberattack and just "ran with it." The attention resulted in a four-month FBI investigation, which Braithwaite said reached the same conclusion that employee error was to blame.

The upside of this, is that all municipal water treatment plants likely **did** and **do** still need to be more alert and aware of threats and to keep their security tight. So if the report of this dastardly attack served to get some inattentive management to tighten up their ship, change employee discharge policies, perhaps delete a bunch of old and unused access credentials, and so forth, then the bizarre inflation of this minor event probably helped security overall.

AI Overlord Hysteria

It's not going to be news to anyone that the sudden explosive popularity, adoption and use of Large Language Neural Network Models,such as ChatGPT, has caught pretty much everyone flat footed. Everyone is running around trying to figure out what it means, whether it's some sort of apocalypse that the public should be protected from.

A couple of weeks ago, the U.S. Chamber of Commerce issued a report from the "AI Commission." (Apparently we have an AI Commission now.) The report claims to highlight the promise of AI while calling for a "Risk-Based, Regulatory Framework." The subhead of the report says: *"Report Finds Policymakers Must Enforce Existing Laws and Develop Policies to Steer the Growth of Responsible, Ethical AI."*

What's actually happening is clear to see: The politicians and bureaucrats have no idea what any of this is, and that scares the crap out of them. And these people are not the brightest bulbs. As our picture-of-the-week explained, it was only recently that a majority of the U.S. Congress understood that TicTok is not a breath mint. Yet these are the people who are going to shepherd us to the land of ethical AI.

Here's what the U.S. Chamber of Commerce is thinking. They wrote:

> *The use of artificial intelligence (AI) is expanding rapidly. These technological breakthroughs present both opportunity and potential peril. AI technology offers great hope for increasing economic opportunity, boosting incomes, speeding life science research at reduced costs, and simplifying the lives of consumers. With so much potential for innovation, organizations investing in AI-oriented practices are already ramping up initiatives that boost productivity to remain competitive.*
>
> *Like most disruptive technologies, these investments can both create and displace jobs.* **[Oh oh!!]** *If appropriate and reasonable protections are not put in place, AI could adversely affect privacy and personal liberties or promote bias. Policymakers must debate and resolve the questions arising from these opportunities and concerns to ensure that AI is used responsibly and ethically.*
>
> *This debate must answer several core questions: What is the government's role in promoting the kinds of innovation that allow for learning and adaptation while leveraging core strengths of the American economy in innovation and product development? How might policymakers balance competing interests associated with AI—those of economic, societal, and quality-of-life improvements—against privacy concerns, workforce disruption, and built-in-biases associated with algorithmic decision-making? And how can Washington establish a policy and regulatory environment that will help ensure continued U.S. global AI leadership while navigating its own course between increasing regulations from Europe and competition from China's broad-based adoption of AI?*

Oh, that's right! You hog-tie U.S. innovators with needless fuzzy and ill-conceived regulations and the rest of the planet is going to leapfrog the U.S with their high speed plans for unethical AI!

> *The United States faces stiff competition from China in AI development. This competition is so fierce that it is unclear which nation will emerge as the global leader, raising significant security concerns for the United States and its allies. Another critical factor that will affect the path forward in the development of AI policy making is how nations historically consider important values, such as personal liberty, free speech, and privacy.*
>
> *To maintain its competitive advantage, the United States, and like-minded jurisdictions, such as the European Union, need to reach agreement to resolve key legal challenges that currently*

> *impede industry growth. At this time, it is unclear if these important allies will collaborate on establishing a common set of rules to address these legal issues or if a more competitive—and potentially damaging—legal environment will emerge internationally.*
> *AI has the capacity to transform our economy, how individuals live and work, and how nations interact with each other. Managing the potential negative impacts of this transition should be at the center of public policy. There is a growing sense that we have a short window of opportunity to address key risks while maximizing the enormous potential benefits of AI.*

Wow. I doubt that anyone knows what any of that really means, yet... but it's clear that we're in for an interesting next five years or so... because it does appear that whatever this is, it's coming on like a speed train... and it's unclear to me that there's anyway to slow it down. It is what it is, and now the entire planet has had a glimpse of a new and uncertain future.

Italy says NO to ChatGPT
I should mention that at least one country has chosen to adopt a much more knee-jerk position on AI. Believe it or not, Italy has banned ChatGPT. I'm not kidding. The news is that the Italian Data Protection Authority has issued a temporary ban on ChatGPT as the agency investigates a possible breach of GDPR regulations, with the agency accusing the OpenAI service of *"unlawful collection of data."*

Now that appears to be the point of this. Someone has apparently freaked out Italian legislators because they believe that "The ChatGPT" is sucking up all of the Internet's data on Italian citizens without any regard for their GDPR-guaranteed privacy.

The only thing I could find to bring some clarity to this was a press release originally written in Italian and poorly translated into English. (The translation would be great if they'd just let ChatGPT do it! But noooooo….) Here's what was announced, from which you can get a sense for their somewhat hysterical concern. Again, pardon the translation, which reads:

> *No way for ChatGPT to continue processing data in breach of privacy laws. The Italian SA imposed an immediate temporary limitation on the processing of Italian users' data by OpenAI, the US-based company developing and managing the platform. An inquiry into the facts of the case was initiated as well.*
>
> *A data breach affecting ChatGPT users' conversations and information on payments by subscribers to the service had been reported on 20 March. ChatGPT is the best known among relational AI platforms that are capable to emulate and elaborate human conversations.*
>
> *In its order, the Italian SA highlights that no information is provided to users and data subjects whose data are collected by Open AI; more importantly, there appears to be no legal basis underpinning the massive collection and processing of personal data in order to 'train' the algorithms on which the platform relies.*
>
> *As confirmed by the tests carried out so far, the information made available by ChatGPT does not always match factual circumstances, so that inaccurate personal data are processed.*

> *Finally, the Italian SA emphasizes in its order that the lack of whatever age verification mechanism exposes children to receiving responses that are absolutely inappropriate to their age and awareness, even though the service is allegedly addressed to users aged above 13 according to OpenAI's terms of service.*

So this is a bit weird, right? So far as anyone knows, no private data has been or is being harvested by ChatGPT. The system is simply ingesting vast quantities of publicly available information and merging it into it an evolving large language model which has an interactive conversational interface. In a very similar fashion, Google's "GoogleBots" crawl the web, following links and indexing everything they encounter. So there's really no difference between what Google collects and what ChatGPT collects. The difference is in the accessibility and presentation of the information. Google maintinas a massive index whereas ChatGPT and similar systems maintain a massive model.

Anyway... I imagine that things will likely calm down in Italy once someone who knows something talks to those who need to know more than they apparently do. One thing seems clear, which is that Italy would not want ChatGPT to be ignorant of everything that country has to offer. It would represent a huge blind spot that would hurt them economically as the world's inevitable adoption of and reliance upon this emerging technology continues. If someone asks their favorite AI-Bot for recommendations about where to stay in Rome, you don't want the world's AI's asking "What's a Rome?"

Twitter code on GitHub

On Friday, Twitter officially posted the source code it uses for selecting which Tweets users will see. This has, of course, been the source of much controversy since the Tweets people see turns out to have an outsized effect upon the beliefs they hold. So, for anyone who may be interested, it's there. Since I'm not a Twitter surfer, I'm not very curious, but I'm sure we'll be seeing some analysis of the algorithm once those who do care have invested the time in learning what it all does. Ars Technica shared their take after a quick look and from what they saw it mostly looked what any student of social science would predict: Watch what each people does, what they click on, what they search for, how long they remain, how far down they scroll, and so forth. Collect every possibly meaningful scrap of data, figure out what sorts of things they want to see, then send them more of that.

One observer did note that Twitter is actively burying Tweets about Russia's invasion of Ukraine. So that mystery is resolved for anyone who might have wondered where those Tweets went after Twitter became an Elon property. We talked about how Elon's Starlink satellite system was turning out to be hugely instrumental in helping to keep Ukraine connected to the rest of the world as Russian missiles continued to batter Ukraine's communications and other utility infrastructures.

It's illegal... How much will that be?

This bit of news struck me as so odd. The government of India is reportedly seeking bids from as many as a dozen me-too Smartphone surveillance software purveyors because the current preferred goto spyware, Pegasys from Israeli's NSO Group's, has become too well known. The use of Pegasys is being frowned on by several annoyed nations, including the current U.S administration, which has been quite vocal. So now it appears to matter which super-secret hidden illegitimate surveillance spyware a country chooses to implant into the smartphones of its surveillance targets.

The news coverage of this in the Financial Times stated that India is hunting for new spyware with a lower profile than the controversial Pegasus system which has been blacklisted by the US government. Indian defense and intelligence officials have decided to acquire spyware from less well known and less publicly exposed competitors of the NSO Group, and they plan to spend around $120 million for the replacement software they feel they need. These lesser known alternatives are doubtless gleeful that Pegasys became so popular that it's now being frowned upon by those with a reputation to maintain. Again, this all seems so bizarre since all of this is completely illegal.

The Financial Times noted that India's move shows how demand for sophisticated, unregulated and illegal technology remains strong despite growing evidence that governments worldwide have abused spyware by targeting dissidents and critics. India, for their part, has never publicly acknowledged ever being a customer of NSO. However, Pegasys spyware has been found on the phones of secular journalists, left-leaning academics and opposition leaders around India, and this sparked a political crisis. As we know, Pegasus turns phones into surveillance devices which can collect encrypted WhatsApp and Signal messages surreptitiously. This is what the spooks inside governments want, claim they need, and are willing to pay for.

And I loved this: India's Modi government officials have grown concerned about the "PR problem" caused by the ability of human rights groups to forensically trace Pegasus, as well as warnings from Apple and WhatsApp to those who have been targeted, according to two people familiar with the discussions. So, yeah... we're looking for something that is just as effective but is also much less well known. Where do we wire the money?

Meanwhile...

Meanwhile, with timing that you couldn't make up, last Wednesday the so-called Summit for Democracy was held, after which a joint 12-nation statement condemning the proliferation of exactly the sort of commercial spyware that India is currently on the mart for. The governments of, in alphabetical order, Australia, Canada, Costa Rica, Denmark, France, New Zealand, Norway, Sweden, Switzerland, the UK, and the US published a joint statement on planned efforts to counter the proliferation of commercial spyware. Agreed countermeasures include tightening export controls, better information sharing to track the proliferation of such tools, and engaging with additional partner governments to reform and align policies on the use of spyware by their agencies. Given how widespread the practice has become at the highest levels of nation state intelligence tradecraft, you really wonder how many of the signatories of that letter are condemning one moment and deploying exactly such spyware the next.

The U.S. FDA & medical device security

Guidance issued by the The US Food and Drug Administration (FDA) last Thursday on March

30th, explains that any and all medical devices being submitted to the FDA for approval will from now on be required to meet specific cybersecurity requirements. These new requirements are part of the Consolidated Appropriations Act which was signed into law in late 2022. That new law contained a section titled "Ensuring Cybersecurity of Medical Devices."

According to the FDA, submissions for new medical devices will need to include specific cybersecurity-related information, such as the description of a plan for identifying and addressing vulnerabilities and exploits in a reasonable time. Companies must also provide details on the processes and procedures for releasing postmarket updates and patches that address security issues, including through regular updates and out-of-band patches in the case of critical vulnerabilities. The information provided to the FDA must also include a software bill of materials (SBOM) for commercial, open source and off-the-shelf components.

The requirements apply to any "cyber device" that runs software, has the ability to connect to the internet, and would be vulnerable to cyber threats. The new cybersecurity requirements do not apply to submissions prior to March 29, 2023, and the FDA will not reject applications solely on this requirement until October 1 — it will provide assistance to companies until that date. However, starting with October 1, the agency may start rejecting premarket submissions that do not contain the required information.

So this is VERY cool and very welcome. The stories about insulin pumps being hacked and taken over have almost become an security industry meme. So this welcome legislation means that devices can no longer be sold and forgotten.

This works easily today for medical devices, the marketing and sales of which are already highly regulated. But it seems likely that somewhere in the future anything that can connect to the Internet may need similar mandatory functionality. It would always be better if it was done voluntarily by manufacturers and if consumers were to insist upon those features.

Hack the Pentagon

Seven years ago, in 2016, we covered the U.S Department of Defense's first, halting, experimental "Hack the Pentagon" Bug Bounty effort. No one was sure it would work – least of all the skeptical government bureaucrats – but work it did. Since that program's launch, ethical hackers have helped the DoD find and fix more than 2,100 vulnerabilities which have been identified by more than 1,400 hackers. The news today is that the continually useful and successful program now has a dedicated website at www.hackthepentagon.mil which the DoD will be using to educate additional branches of the sprawling U.S government about the benefits available from allowing good-guy hackers to take a crack at cracking. The site will also continue to seek and sign-up talented hackers for the continually expanding program.

Miscellany

Firefox 3dr-party DLL check-up

Since Firefox 110, which everyone should have by now (I'm at 111.0.1) Firefox has provided a built-in facility for showing 3rd-party DLLs, meaning not signed by either Mozilla or Microsoft, which have arranged to have themselves injected intoFirefox's address space. If you're a Firefox user, type “about:third-party” to view this new page. When I did that I only had three DLLs signed by Intel. Since I’m running on an Intel NUC machine that wasn’t surprising. By clicking on the little folder icon, Windows Explorer will open on that file and it’s possible to then right-click and check out its properties. And sure enough, in my case the three DLLs were Intel graphics drivers signed by Intel.

This is useful from a security standpoint, since injecting DLLs into browser process space is something that malware likes to do to steal things like passwords and crypto addresses. But Mozilla’s primary motivation was to help identify misbehaving DLLs that might be causing Firefox to become unstable and to crash. Next to the little folder icon there may be another toggle that allows the user to prevent a module’s injection.

Closing the Loop

Microsoft’s Extortion?

Not surprisingly, a great many of our listeners wrote after last week’s Microsoft rant. Rather than share with everyone here what was essentially the same sentiment expressed over and over, I’ll just share one that’s representative of all:

Matthew Hile / @mhile

> *Steve,*
> *Listener since episode 0001 and SR alpha tester (thanks!). I think you need to reconsider last weeks rant. For years you have decried outdated unpatched servers stuck in the closet. You have spoken positively of governments' recent efforts to locate and report vulnerable servers. You have asked what it would take to get those servers updated and would it ever happen. Well MS's refusal to accept email from unpatched and unsupported exchange servers is clearly a very dramatic way to accomplish your goal. Regardless of the "dangers" from email from those systems this will force folks to either upgrade or move from known vulnerable systems to a less vulnerable one making the internet safer for all.*

I agree **completely** with what Matthew wrote, and with what everyone who wrote something similar, said – 100%. Really. This is an absolutely powerful and doubtless effective means for Microsoft to force the upgrading of their older Exchange Servers. And toward the end of last week’s rant I even said exactly that, though I would not fault anyone for missing that, since it certainly wasn’t my main thrust. What I said toward the end of my rant was:

“No one using any Microsoft Exchange Server software will ever again be able to fail to keep it updated, nor to avoid the purchase of future licenses – forever. ***I celebrate that idea from here forward, since keeping software updated, especially Exchange Server, is a good thing…****”*

So, I just wanted to make sure that everyone understood that I really do appreciate that aspect of this mess. But in my mind, the fact that it **would** be effective doesn’t make it right.

Sci-Fi

The Silver Ships

It was episode #887, which we recorded on September 6th of last year. Before I went back to look it up, I assumed that it was longer ago, because that means that it's only been 7 months since I shared my discovery of this new-to-me author. And in that time I've read his first two series totaling 24 books. Our listeners will know that I'm referring to The Silver Ships series by Scott Jucha.

Leo was not a fan of the series, since he felt that the main character, Alex Racine, was unrealistically designed to be too perfect. I understand that. Not everyone is going to like everything. Not everyone likes the same sort of music or the same food. But, for what it's worth, I had a wonderful time with the series and I've heard from many of our listeners who felt the same way. And I mean, rave reviews. There was some terrific science fiction in there and Scott Jucha is a great story teller.

When I look back over the 24 books, my main complaint would be that 20 books spent with the same characters is a lot of time with them. I knew that was the case because when I switched to the 4-book Pyreans sideline I was a bit relieved to be meeting some completely new characters in an entirely different set of worlds. And, inevitably, any 20-book storyline is going to have some spots where you're wading through detail that just seems to be taking up time.

I'm mentioning the series again because I did love it and I'm glad to have recommended it to everyone here. If it wasn't your cup of tea, then no harm done. And if it was, then you already know how much fun was contained within those pages.

Since the beginning of the year, the amount of time I've allowed myself for recreational reading has been significantly curtailed, because all I really want to do is get SpinRite v6.1 finished and published. So whenever I'm awake, except for time spent assembling this podcast each week, SpinRite has had my attention. But I did manage to finish the final 20th book of the Silver Ships series.

Since I really think that Jucha is extremely pleasant to read, I have opened the first of his 8-book "Gate Ghosts" series. And so far I love it too. New people, new worlds, but the same very acceptable writing and terrific storytelling.

And I should also note that while I've been doing this, another favorite author of ours, Ryk Brown, has been cranking away on the third one of his 15-book arcs as he continues to tell the story of Nathan, Jessica, Cameron, Telles, Josh, Loki and the rest. Anyone who has dipped their toe into the Frontiers Saga knows all of those names quite well. Once I eventually finish Jucha's 8-book Gate Ghosts novels, I'll switch back to catch up with the ongoing Frontiers Saga.

SpinRite

And speaking of SpinRite, very briefly, we’re getting very very close. SpinRite hasn’t misbehaved in a long time since it adopted full protection from any of the many ways that BIOS firmware can misbehave. And there are many. But I think that if anything it’s probably over-insulated. The few problems remaining mostly surround how hard SpinRite should try to work with badly damaged drives. When a drive is really very badly damaged it’s even difficult for SpinRite to verify that it has established communication with the drive. So we’re sorting through that now. On one hand it’s mostly of academic interest since none of the clearly dead and dying drives that SpinRite’s testers have would ever be considered useful for data storage. But in wanting to make it as good as it can be, and in not wanting to ever need to return to it once it’s declared finished, I’m still willing to give it a little more time... But not much more.

Zombie Software

We know that a few years ago CISA created their KEV which is the Known Exploited Vulnerabilities list or database. The idea was for this list to serve as a prioritization for any entity exposed to the Internet since it would not reflect ever vulnerability ever seen. No. CISA's KEV is specifically for vulnerabilities whose active exploitation has recently been observed in the wild. In other words, they are active vulnerabilities.

We also talked about how the size of CISA's KEV ballooned last year, but that was not because of many new vulnerabilities discovered with patches available in 2022, but rather because older vulnerabilities — in some cases almost ancient — were seen still in use and were therefore, as is KEV's charter, added to the list.

Now, today, we have another shoe dropping, with a rather breathtaking report from a security firm known as Rezilion. Here's what Rezilion wrote:

> *Do you know KEV? You should, because hackers do! Rezilion's research team just released a new report, which highlights the critical importance of Known Exploited Vulnerabilities (KEV). Specifically, our research uncovers that although KEV catalog vulnerabilities are frequent targets of APT Groups, many organizations are still exposed and at risk from these vulnerabilities because they are not patching them. This gap in patching may be due to a lack of awareness, or a lack of patching resources, or both.*
>
> *The KEV catalog, maintained by the Cybersecurity and Infrastructure Security Agency (CISA), is a reliable source of information on vulnerabilities that have been exploited in the past or are currently under active exploitation by attackers. According to our new research, there are over 15 million vulnerable instances in the KEV catalog, with the majority being vulnerable Microsoft Windows instances. That's a massive number of systems exposed to attacks, leaving organizations vulnerable to exploitation from threat actors and Advanced Persistent Threat (APT) groups.*

So, Rezilion scanned the Internet checking specifically for systems exposing vulnerabilities to any of the 896 vulnerabilities currently listed in the CISA KEV database. What they found was what any bad guys who took the time to do the same could find, which was more than **15 million** systems worldwide currently in need of patching to prevent their easy exploitation. Rezilion's report listed the top ten most frequently exposed vulnerabilities in a table which I've included in the show notes (see next page).

I was curious to know more about the specifics of a few of these top 10 still-most-prevalent exposed vulnerabilities. So I looked them up.

#1 on the top 10, currently present at 6,453,785 IP addresses, is CVE-2021-40438. So that one is only two years old.

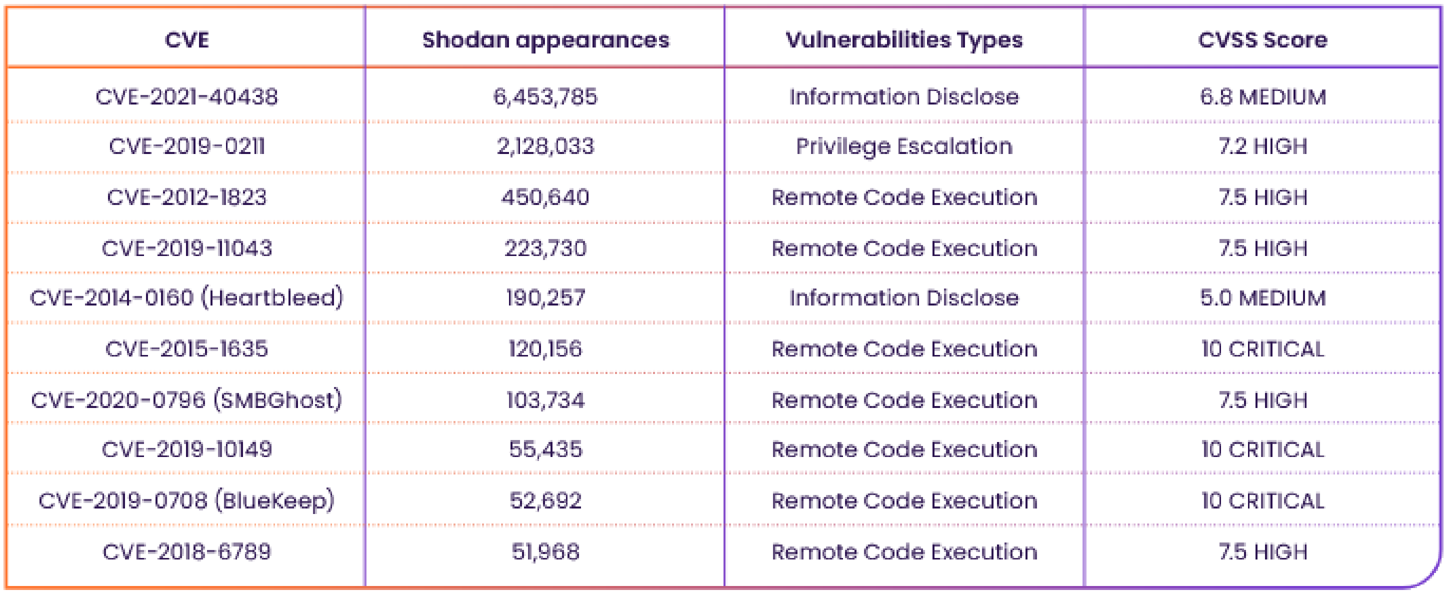

CVE	Shodan appearances	Vulnerabilities Types	CVSS Score
CVE-2021-40438	6,453,785	Information Disclose	6.8 MEDIUM
CVE-2019-0211	2,128,033	Privilege Escalation	7.2 HIGH
CVE-2012-1823	450,640	Remote Code Execution	7.5 HIGH
CVE-2019-11043	223,730	Remote Code Execution	7.5 HIGH
CVE-2014-0160 (Heartbleed)	190,257	Information Disclose	5.0 MEDIUM
CVE-2015-1635	120,156	Remote Code Execution	10 CRITICAL
CVE-2020-0796 (SMBGhost)	103,734	Remote Code Execution	7.5 HIGH
CVE-2019-10149	55,435	Remote Code Execution	10 CRITICAL
CVE-2019-0708 (BlueKeep)	52,692	Remote Code Execution	10 CRITICAL
CVE-2018-6789	51,968	Remote Code Execution	7.5 HIGH

Here's what Rapid7 wrote about this back on November 30th of 2021:

> *On September 16, 2021, Apache released version 2.4.49 of HTTP Server, which included a fix for CVE-2021-40438, a **critical** server-side request forgery (SSRF) vulnerability affecting Apache HTTP Server 2.4.48 and earlier versions. The vulnerability resides in mod_proxy and allows remote, unauthenticated attackers to force vulnerable HTTP servers to forward requests to arbitrary servers — giving them the ability to obtain or tamper with resources that would potentially otherwise be unavailable to them.*
>
> *Since other vendors bundle HTTP Server in their products, we expect to see a continued trickle of downstream advisories as third-party software producers update their dependencies. Cisco, for example, has more than 20 products they are investigating as potentially affected by CVE-2021-40438, including a number of network infrastructure solutions and security boundary devices. To be exploitable, CVE-2021-40438 requires that mod_proxy be enabled. **It carries a CVSSv3 score of 9.0.***
>
> *Several sources have confirmed that they have seen exploit attempts of CVE-2021-40438 in the wild. As of November 30, 2021, there is no evidence yet of widespread attacks, but given httpd's prevalence and typical exposure levels (and the fact that it's commonly bundled across a wide ecosystem of products), it's likely exploitation will continue — and potentially increase.*

This vulnerability would not be leveraged for any sort of widespread attack. But it is exactly what a high-level determined and targeted attacker might be looking for. Imagine a publicly exposed Apache web server on an enterprise's network boundary, as most are. The web server sits on the enterprise's network, as it pretty much has to, and it fields public HTTP requests over port 443. And that the enterprise also maintains other servers of various types for strictly internal use and those servers have no public exposure; they are simply sitting on the internal network. However, they are visible to the vulnerable Apache web server since they reside on the same internal corporate network. This vulnerability presents the perfect means for allowing a remotely located attacker to bounce queries off of the enterprise's public server and into the private corporate network behind it. It turns the Apache server into an unregulated and unrestricted public proxy with access to the enterprise's internal private network. If and when used in that way, it's a horrifying vulnerability... and, right now, of the 896 known vulnerabilities on CISA's KEV list, it holds the #1 slot by being present at 6,453,785 known IPs spread around the world.

The presence of this vulnerability came to light only about 18 months ago. And one of the problems, as noted by Rapid7, is that Apache is embedded in many networked appliances. They noted that Cisco has at least 20 different devices that were and, probably to a distressing degree still are, subject to this vulnerability. So it's not just instances where someone is running a standalone Apache web server on a Unix or Linux box. Those would be obvious and easily updated. No. The real challenge is all of the places where Apache has been embedded inside an appliance that appears to be humming along just fine. "Cisco has some updates for our network thing-a-ma-jig? Well, okay, but it's working just fine now. We'll get back to that when we have some time." — Then suddenly your entire network is encrypted, the ransom note arrives, and you **really** don't have any time. You wonder after the fact how they got in? They pivoted off of a known but unpatched vulnerability, 6,453,785 instances of which currently exist and were recently enumerated on the public Internet.

Now, I'm not going to go through all 10 of these, but let's look closely at one that's halfway down the list occupying the #5 slot. That's CVE-2015-1635. The first thing we notice is the 2015 – so, eight years ago. And it was in April of that year, so it's been a full eight years. You might also notice the quaint 4-digit 1635 number of the CVE from that year – ahhhhh, simpler times. And how many known vulnerable instances of that now, 8-year old problem, are still presently exposed on the public Internet? 120,156. And get this, it's one of those quite rare vulnerabilities that earned itself a CVSS score of 10.0. That's right. It doesn't get any worse or more frightening for those who rate these things. We've seen how difficult it is to score a perfect 10. There are plenty of wannabe 9.8's out there, but the 10.0 remains a rare beast. And who managed to bring that one home? None other than Microsoft in their own web server where its successful exploitation – simply by making the proper remote query – allows the attacker's code to run in the user's system with full SYSTEM kernel level privileges – thus the full monty 10.0 score. Again, a full 8 years downstream and yet today 120,156 Microsoft web servers remain sitting ducks for this remote code execution vulnerability.

I'm a software developer who has done a fair share of Internet programming. SQRL is an Internet authentication system, GRC's ShieldsUP and DNS Spoofability services are persistently on the Internet, and GRC's DNS Benchmark with 8.2 million downloads, now occurring at a rate of around two thousand per day, has become the industry standard for measuring DNS server performance. I published it 13 years ago, back in 2010, and it's never had a bug.

So I deeply and fundamentally object to what is clearly a growing presumption in the world that any software that's not being actively maintained is therefore inherently bug ridden. It's as if no one trusts software anymore that isn't **in** intensive care with multiple IV bags hanging overhead. It has to be in serious trouble, on continuous life support, being monitored 24 hours a day for anyone to think that it **doesn't** have serious problems. It's perverse.

And yet, the most frustrating thing is, I can't argue against that, because it **is** the only sane and rational conclusion to draw from all of the evidence that is presented to us everyday.

There are some beautiful exceptions to that. SpinRite 7 will be based upon an embedded 32-bit real time operating system, all ongoing support for which **ended** at the end of last year when its publisher went out of business; and yet I've selected it for SpinRite's future. Why? They didn't go out of business because the software was no good. They went out of business because the software was too good. It was done. It was a perfectly functioning finished product and no one was subscribing to updates anymore because it was feature complete and there were no more bugs to be fixed. None. It was done. Take it off the ventilator. Pull the IV lines. Hold your breath and count to 10. Does it still have a pulse? What do you know! It's alive! Works great.

Now, the idea that a company went out of **business** because its software **didn't** have any bugs is a bit chilling because **nobody** wants to use Windows after its life support has been terminated. If it's still moving after that, the presumption is that it's become a zombie that wants to eat your brain. So when we hear that Microsoft shipped a brand new edition of Windows – I don't recall which Windows it was – Vista, 7, 8 or 10 – which, at the time of its release had more than 10,000 bugs known to Microsoft, you have to wonder whether the leaking, of what should have been embarrassing information, might not have been deliberate on their part. If a company went out of business because their software worked perfectly and had no more bugs, that's not a worry that will ever keep anyone at Microsoft awake at night.

Everyone wants to use Windows, and everyone wants to make sure that Windows' IV lines are continuously connected and that vital, life sustaining fluids, are flowing freely. And the minute Microsoft cuts off that flow, **whoa!** ... time to upgrade. And not because you necessarily want anything that the new Windows has. Quite often, in fact, much of what they have done to it is noxious and unwanted – just look at Vista or Windows 8. But what choice do you have? Everyone knows that Windows is perpetually so sick that the only safe way to use it is with it lying in ICU, attached to a ventilator, with IV lines flowing and under constant monitoring. You disconnect it from its continuous life support at your peril.

And Microsoft with Windows, or Exchange Server, or, well, I guess anything Microsoft creates, is only a slice of a much larger problem. This is the sad reality of today's software ecosystem. Our insatiable desire for features, or in Microsoft's case their unrelenting need to seduce us with new bug-ridden features, means that we're never going to get ahead of this, and that, indeed, almost all software only remains viable while it's on perpetual life support.

So now we loop back to CISA's KEV list, and we see that what we have is a list of what can only be described as Zombie Software. For one reason or another, that software has been disconnected from the life support that was, sadly and unnecessarily but still, in fact, vital for its ongoing safe use on the Internet.

In the case of that eight year old Windows web server critical 10.0 vulnerability, that Windows 7 and Server 2008 era software is now more than three years past its end of support life, which Microsoft terminated in January of 2020. It's worth noting, however, that the updates for those 120,156 very old servers are still online and are available from Microsoft. So, simply clicking "check for updates" on any of those machines, assuming that they're activated and validly licensed, will quickly bring them current with the last build of their server's code and would, in the process, cure any and all vulnerabilities that were known and for which patches were available when support for that server ended.
But that's one of the two oldest entries on the list. The other eight are much more recent and

ongoing real-time life support still remains available for all of the others. A perfect example of that is the other critical vulnerability we examined in Apache. ANY of those 18-month old Apache servers could be updated immediately. Yet 6,453,785 publicly exposed instances of Apache **voluntarily** remain unpatched today, despite the fact that updates are certainly readily available for every instance of them.

So where does this leave us?

One thing I want to observe is that zombies eat the brains of those who are nearby. So this is about self responsibility. If you don't want your zombie server to eat your brain, either shut it down or get it patched, if possible. This is one of my main complaints with Microsoft's plan to force **other** people to update their Zombie Exchange Servers. Microsoft contends that **their** brain was going to be eaten by **someone else's** remotely located zombie. That's utter nonsense. We all understand that we're the ones whose brains are in danger if we choose to use software after it's gone zombie.

But taking the broadest view possible, it is an unfortunate and sad truth that with very few exceptions, almost none of today's software is able to stand alone without the presumption that periodic and continuing updates are required for that software's users to feel comfortable with the performance of their software.

The assumption of the availability of the ever present Internet, which provides a life line for would-be zombies, allowing them to remain on life support, has enabled a gradually growing laxity on the part of software publishers. Today, few apparently feel the need to get it right the first time since problems can always be fixed later.

In the extreme case of Microsoft, major products are shipped, despite reportedly being riddled with tens of thousands of bugs known to its developers at the time, and as we saw with last week's Exchange Server issue, Microsoft's own self-fulfilling prophecy is that once any of their software is removed from life support, not only everyone nearby, but also everyone else on the Earth, is put in danger once it becomes a Zombie.

It's a sad state of affairs which didn't have to be. But, as we sometimes observe on this podcast, it's the world we're in. The best we can do is to navigate it safely.