Orange
Cyberdefense

Security Intelligence | Research Report

Monthly Report
February 23

Contents

Introduction

The pledge made by some ransomware gangs not to target healthcare organisations and other medical facilities during the pandemic is now very much a distant memory.

The healthcare industry is now firmly back on the table as a target for most of the main Cy-X groups, with Clop, Royal, BlackCat & KillNet, as well as some other state sponsored groups, all known to target the healthcare sector.

Two recent examples serve to highlight this point. Firstly, in what can only be described as a new low the BlackCat group, also known as ALPHV, published photographs of breast cancer patients it had stolen from the US based Lehigh Valley Health Network in an attempt to extort a ransom payment. In a second incident the non-profit organisation Methodist Family Health, who describe themselves as "a complete continuum of care for Arkansas children who are abandoned, abused, neglected and struggling with psychiatric, behavioral, emotional and spiritual issues." were targeted by the AvosLocker group. According to the AvosLocker leak site the exfiltrated data includes thousands of patient EMR (Name, Address, Social Security Number, Date Of Birth, Medications), accounting and payroll files.

At a glance

The healthcare industry is now firmly back on the table as a target for the majority of the main Cy-X groups.

World Watch Review

The Orange Cyberdefense CERT published a total of 16 new World Watch advisories during February 2023, along with adding updates to a further 18 previously published advisories.

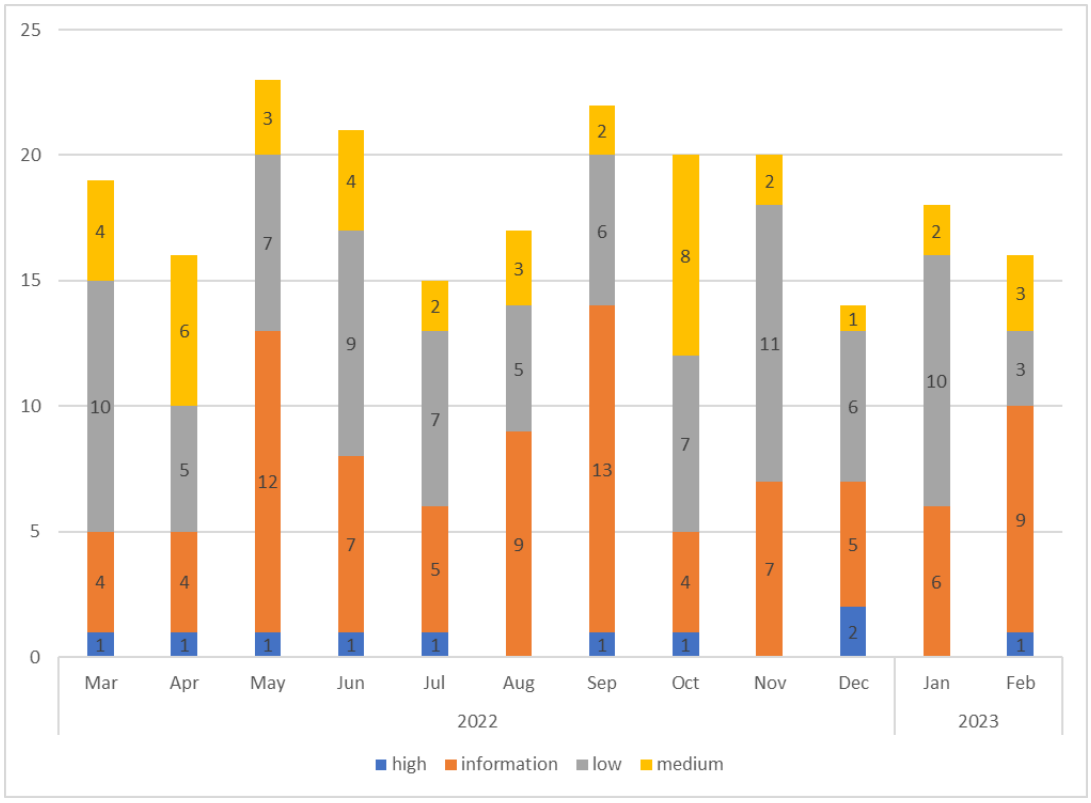

Breakdown of Published Advisories Previous 12 Months

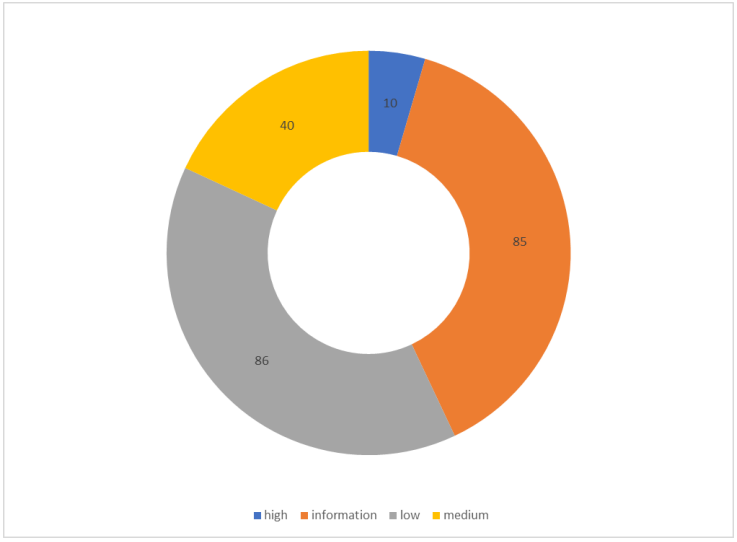

Breakdown of Advisory Criticality for Previous 12 Months

Advisory Summary

As can be seen above we had one high criticality advisory this month with the remaining 15 advisories given criticality ratings of low, medium or information when initially published. These ratings are based on our CERT's assessment of the risk and threat levels associated with the subject of the advisory at the time of publication, so even though an advisory may concern a vulnerability rated as critical by the vendor we may deem it to only initially be medium, if say there is no publicly available exploit. This is under constant monitoring however and subsequent updates will increase our criticality level as required if circumstances should change. Some advisories of note this month are:

678367 - New ESXiArgs ransomware campaign targets ESXi servers

- Since yesterday, a massive campaign by a possibly new ransomware threat actor impacts many outdated ESXi servers throughout the world. Indeed, several of our clients as well as at least 3 cloud providers hosting such servers for their clients warned us privately that some servers running old versions of ESXi had been encrypted. Servers running versions 6.x and maybe 5.x, that are currently End-of-Life thus not fully maintained by VMware, are in the scope of this attack.

- One 2-years old vulnerability in ESXi's OpenSLP, tracked as CVE-2021-21974, has been most presumably leveraged at least in some of these incidents. Patched in February 2021 by VMware, this vulnerability was not known to be massively exploited in the wild. A PoC was nevertheless publicly disclosed soon after the fix was released, as well as a stable exploit. Another less likely vulnerability could be the one VMware disclosed last December, as Scaleway (a French hoster) mentionned publicly in a tweet. The patch includes in particular a vulnerability numbered CVE-2022-31696, that enables an attacker to escape a sandbox and thus impacts the full ESXi servers' VMs (i.e. multiple clients). Nonetheless, this issue requires local access.

682867 - PoC now available for a critical RCE vulnerability in Microsoft Word

- On March 5, security researcher Joshua J. Drake released a PoC for the CVE-2023-21716. Also fixed during Microsoft February Patch Tuesday, this heap corruption vulnerability located in Microsoft Word allows an unauthenticated remote attacker to execute arbitrary code using a specially forged RTF file. Unfortunately, according to Microsoft, previewing the file with "Preview Pane" is enough to exploit this vulnerability. Our experts at Vulnerability Intelligence Watch have attributed it a CVSS score of 8.5 out of 10.

- This vulnerability impacts the following versions of Microsoft Office:

    - Microsoft Office 365
    - Microsoft Office 2016
    - Microsoft Office 2013
    - Microsoft Office 2010

- The PoC provided by the researcher is a python script that allows a user to generate a file that triggers the security issue. Nevertheless, this script does not generate an exploit, so threat actors must create one themselves in order to exploit this vulnerability.

684035 - PoC released for 0day vulnerability in Windows Contacts

- Security researcher j00sean recently published a technical write-up, including proof of concept code, detailing how to exploit CVE-2022-44666 to achieve remote code execution on up-to-date Windows systems. The vulnerability was first discovered and reported to Microsoft back in 2018. At the time, the vendor decided not to fix it, so it was disclosed as a 0-day by ZDI on January 2019. Microsoft finally addressed the vulnerability with the December 2022 monthly security update, but the fix is incomplete and easily bypassable, as shown in the report.

- The vulnerability affects the Windows Contacts application and could be triggered by opening a .contact or .vcf file and then clicking on a link, which could be displayed as the contact's email address or website. Thus, user interaction is required for this exploit to work. The vulnerability received a CVSSv3 base score of 8.8 from our Vulnerability Intelligence experts.

- Finally, 0patch has released a free unofficial patch for all affected Windows systems.

684001 - Critical vulnerability fixed in the latest version of Joomla

- On February 13, the Joomla Security Strike team announced an important security patch for their product, a famous CMS among individuals and small organizations. Numbered 4.2.8, this latest version fixes a critical vulnerability present in all versions between 4.0.0 (released on August 17, 2021) and 4.2.7. Tracked as CVE-2023-23752, this vulnerability located in one REST API allows an unauthenticated remote attacker to create a specially crafted request to access webservice endpoints and thus gain access to the global configuration of the application. For example, it appears that it is possible to retrieve the account number and password from the database. That's why Joomla recommends all impacted users to renew various global configuration credentials.

- This vulnerability exists because of inappropriate access restrictions to the web service endpoints. The attacker can bypass the security restrictions in place and presumably compromise the full web application.

- Unfortunately, several PoCs seem to have been made available publicly and appear, if confirmed, to be very trivial to use with a simple GET request. If these PoCs are confirmed, it is very likely that threat actors will exploit the vulnerability to compromise web applications exposed to the Internet.

Editor's Notes

Here the team will provide commentary on a news item, expansion on something specific such as a single incident, or could be as expansive as coverage on trends observed in the threat landscape.

Charl

LastPass lessons learned

If you haven't had a chance to review the various reports and commentary on the recent series of LastPass compromises, then I really recommend you do. The saga started in August last year and has continued to unfold right up to the start of this month.

We've covered the incident and its implications in a series of World Watch advisories that can be accessed here:

https://portal.cert.orangecyberdefense.com/worldwatch/642459

These two links also provide a good summary:

https://arstechnica.com/information-technology/2023/02/lastpass-hackers-infected-employees-home-computer-and-stole-corporate-vault/

https://blog.lastpass.com/2023/03/security-incident-update-recommended-actions/

This incident, which impacted a mature and respected cyber security vendor, is sobering in many respects and warrants close examination with a view to some critical introspection.

There is still a lot about this incident that we don't understand, but for me what stands out is **how comprehensive and intelligent the whole attack seems to have been**. Some commentators have described the attackers as '**smooth'**, and in this context that isn't a good thing.

On August 25 2022, Password management firm LastPass confirmed a breach of their systems that had occurred 2 weeks prior, leading to the exfiltration of some of the company's source code and proprietary technical information.

A long and sometimes hysterical debate ensued about what private information had been exposed, and what had remained properly protected by encryption. This turned out to be mostly a red herring, as the client data stolen was by and large properly protected at the time.

But the LastPass investigation then subsequently revealed that the threat actor had had access to LastPass internal systems during a four-day period in August 2022.

The threat actor behind this attack had managed to gain access to LastPass' development environment by using a developer's compromised endpoint, allowing them to bypass MFA, since the developer had already successfully authenticated using MFA on their endpoint. However, the password manager company argued

that their internal system design and controls prevented the threat actor from accessing the production environment or any customer data. LastPass also analyzed their source code and production builds and confirmed that there was no evidence of attempts of code-poisoning or malicious code injection.

Then, on November 30, LastPass notified their customers about another security breach affecting a third-party cloud storage service shared between LastPass and its parent company GoTo. GoTo also notified their customers in a separate post, adding that they had also detected unusual activity within their development environment.

According to the notice released by LastPass, the threat actor was able to access certain elements of customer information by using information obtained by them during the first intrusion in August 2022, though passwords stored in LastPass apparently remained safely encrypted due to their Zero Knowledge architecture.

The threat actor had copied information from a backup that contained basic customer account information and related metadata, as well as a backups of customer vault data from the encrypted storage container. The customer vault data was encrypted and can only be decrypted with a unique encryption key derived from each user's master password. However, some of the data in the vault such as website URLs is stored in plaintext, which means the threat actor could use it to identify interesting targets based on the domains present.

On January 23, Lastpass owner GoTo updated their blog post about the recent breach for the first time since November 30, when they confirmed "unusual activity" within its development environment and cloud storage service. According to the company, the hackers were able to steal encrypted customer backups from a cloud storage service, alongside an encryption key for "a portion of the encrypted backups".

The stolen data comes from the following GoTo products:

- Central
- Pro
- join.me
- Hamachi
- RemotelyAnywhere
- Rescue
- GoToMyPc

The stolen data includes account usernames, passwords (salted and hashed), a portion of MFA settings, as well as some product settings and licensing information.

GoTo did not release any information on the number of affected customers.

According to a February update by the password management company, the threat actor leveraged information stolen during the first incident, as well as information

available from a third-party data breach along with a RCE vulnerability in a third-party media software package to launch a coordinated second attack targeting the company's encrypted Amazon S3 buckets.

During the first incident, a software engineer's corporate laptop was compromised, giving the threat actor access to a cloud-based development environment. This allowed them to steal source code, technical information and certain LastPass internal system secrets. Then, using information obtained during this incident, the threat actor continued engaging in reconnaissance and enumeration activities.

Then, the attacker targeted a senior DevOps engineer's home computer and exploited a RCE vulnerability in a media software package to implant a keylogger malware, which the threat actor then used to capture the employee's master password for the engineer's LastPass corporate vault (seemingly because the employee accessed their corporate vault from their home computer).

During the second incident, the following data was exfiltrated:

- DevOps secrets, which allowed access to LastPass cloud-based backup storage
- Configuration data
- API secrets
- Customer metadata
- Backups of all customer vault data
- Backup of LastPass MFA/Federation database, containing copies of LastPass Authenticator seeds and telephone numbers used for MFA.

The use of valid credentials made it difficult for the investigators to detect the attacker's activity, allowing the threat actor to access and steal data from LastPass' Amazon S3 buckets for over two months, between August 12, 2022 and October 26, 2022.

The intrusion was finally detected through AWS GuardDuty alerts after the threat actor tried to use Cloud Identity and Access Management (IAM) roles to perform unauthorized activity.

There is (unsubstantiated) speculation now that some technical aspects of the attack point to a North Korean (DPRK) state-backed threat actor, and that the goal might be to retrieve passwords for other high-value accounts, potentially for crypto exchanges. This of course evokes memories of a similarly sophisticated attack against the security business RSA in 2011, in which cryptographic seeds for their secure One Time Pin generators were stolen. See https://www.wired.com/story/the-full-story-of-the-stunning-rsa-hack-can-finally-be-told/ for more.

The broad scope and persistence of this attack should serve as a case study for future penetration testing and red team engagements as well as incident response exercises!

Wicus

Looking at the Exploit Prediction Scoring System

Identifying, prioritizing, and patching vulnerabilities can be a challenge given that the number of vulnerabilities added to the National Vulnerability Database (NVD) is growing each year. A total of 25,068 Common Vulnerability Enumeration (CVE) entries were added to the NVD in 2022, 24.3% more than in 2021. To put this in perspective the number of CVEs published between June 1, 2016, and June 1, 2018, numbered 25,159 vulnerabilities.

The Common Vulnerability Scoring System (CVSS) is a metric to communicate the severity of a CVE and is a static value captured as the Base CVSS. There are two other metrics that augment the Base CVSS metric, namely the Environmental CVSS and the Temporal CVSS. The Environmental CVSS metric is specific to the organization where the vulnerability is present and is adjusted by the perceived risk priorities of the respective entity. The intention of the Temporal CVSS metric is to be adaptable with respect to what is happening in the threat landscape and communicate the availability of exploits for a specific vulnerability, the reliability of these exploits, the remediation options available to reduce impact or eliminate the vulnerability, and the certainty (confidence) that the vulnerability exists. The Temporal CVSS metric is thus more fluid than the Base CVSS, but it is also limited. The Temporal CVSS metric only speaks to potential exploitability and does not speak to active exploitation [1].

System defenders need a way to know how to prioritize actions based on current information in relation to vulnerabilities in their environment. The mere fact that a vulnerability is exploitable does not necessarily require immediate action. Past efforts to communicate the likelihood of the exploitation of vulnerabilities are either proprietary, private, or vendor-specific and thus preclude wide adoption. The Temporal CVSS metric is the closest there is to an open publicly available value, but it's only useful if it is updated.

The Exploit Prediction Scoring System (EPSS)[2] Special Interest Group (SIG) was created to provide a simple public metric that communicates the likelihood of a vulnerability being exploited daily. The EPSS SIG was formed in early 2020 and operates under the auspices of the Forum for Incident Response Teams (FIRST). Version 1 of EPSS, released in 2021, was relatively simple and used 16 features as part of the model to generate a score for each vulnerability on the likelihood of exploitation within the first year of the CVE. The EPSS v1 model could be run locally with little resource requirements. EPSS version 2, released in early 2022, was much more ambitious and switched to a centrally computed model that used 1164 features to calculate an EPSS score for each CVE. One big advantage of version 2 over version 1 was that version 2 could predict exploitation likelihood for the

[1] https://www.balbix.com/insights/temporal-cvss-scores/
[2] https://www.first.org/epss/

following 30-day window instead of for the following year, which is much more practical for defenders.

The latest EPSS version, version 3, was made available to the public on March 7, 2023 and expanded the model to use 1477 features. An important factor for EPSS version 3 was to improve the classifier performance (precision-recall) and version 3 boasts an 82% improvement over version 2.

The EPSS version 3 score for a CVE is calculated using the following features:

- Availability of published exploit code at Exploit-DB, Github, and Metasploit targeting the CVE.
- Inclusion of the CVE in public vulnerability lists such as Google Project Zero's zero-day vulnerability list, Trend Micro's Zero Day Initiative (ZDI) list, and the CISA's Known Exploited Vulnerability (KEV) catalog.
- Mentions of the respective CVE on social media such as Twitter.
- Presence of the CVE in Offensive Security Tools such as Nuclei, Jaeles, Intrigue, and Sn1per.
- The various classes of references that cite the CVE as found under the references section of the MITRE CVE record for example links to Vendor Advisories, News Articles etc.
- Keyword description of the vulnerability that provide attributes of the CVE such as "remote attacker", "code execution", "denial of service", etc.
- The CVSS metrics defined for each CVE[3].
- The CWE ID of the CVE.
- The impacted vendors associated with the CVE as extracted from the CPE value.
- The age of the vulnerability as calculated by the number of days elapsed from CVE publication to model execution.

These sets of parameters are collected for each CVE and make up 1477 variables that are provided as input to an algorithm called XGBoost, a well documented gradient boosted tree algorithm. The output is an EPSS score for each CVE that is an absolute score and can be scaled as part of other calculations.

EPSS version 3 boasts an impressive 82% improvement over the previous EPSS version 2 classifier, but what does this mean for security practitioners? One way to examine the value of EPSS is to consider competing prioritizing strategies. When prioritizing patching vulnerabilities, teams need to determine what they can achieve given the time and budget at their disposal. This induces an inherant conflict between coverage and efficiency4.

'Efficiency' and 'coverage' are terms selected as the more formal terms 'precision' and 'recall' are not intuitive for security practitioners. Low efficiency represents

[3] For CVEs with only CVSS v2 metrics a classifier was created to infer the CVSS v3 metrics
[4] The term efficiency is used by the EPSS SIG in their paper.

wasted effort. Coverage represents how well a remediation strategy is addressing the vulnerabilities.

A small team with limited budget, will seek to be as efficient as possible thus forgoing the option to patch as many CVEs as possible. To illustrate efficiency, one can ask the question: “Given the vulnerabilities that were patched, how many of those were actually exploited?”. If 10 vulnerabilities were patched and 7 were exploited somewhere in the wild, then the efficiency is 70%.

Teams with bigger budgets and more people can afford to favor coverage. To understand this approach, consider the following question: “Given all these actively exploited vulnerabilities, what proportion did the team actually patch?”. In other words, if 10 vulnerabilities were exploited and 5 were remediated the coverage is 50%.

Using previous research by the main authors of the latest EPSS v3 paper, they claim that 5.5% of all vulnerabilities are exploited in the wild [3]. This value is used as the basis for the rest of their calculations.

Using the EPSS scores can give security teams the ability to calculate the level of effort as a percentage of CVEs to patch to reach a desired efficiency or coverage level. The FIRST EPSS SIG shared some figures to illustrate the level of effort (% of CVEs to patch) in terms of four strategies, namely using the CVSS v3 score, EPSS v1 score, EPSS v2 score, or using the EPSS v3 score as a threshold value or minimum score.

When using the CVSS v3 metric and only selecting those CVEs with a CVSS score of 9.1 or higher, results in a 15% level of effort representing 28,000 vulnerabilities that must be patched. The calculated efficiency will be 6.1% and the coverage will be 33.5%.

Comparing EPSS v1 and keeping the level of effort pegged at around 15% will require using an EPSS v1 threshold of 0.062, resulting in a coverage score of 57% and an efficiency of 15.5%. Repeating this for EPSS v2 with a threshold of 0.037 will yield an efficiency of 18.5% and a coverage of 69.9% while keeping the level of effort around 15%. EPSS v3 with a threshold of 0.022 gives us a coverage score of 90.4% and an efficiency of 24.1%.

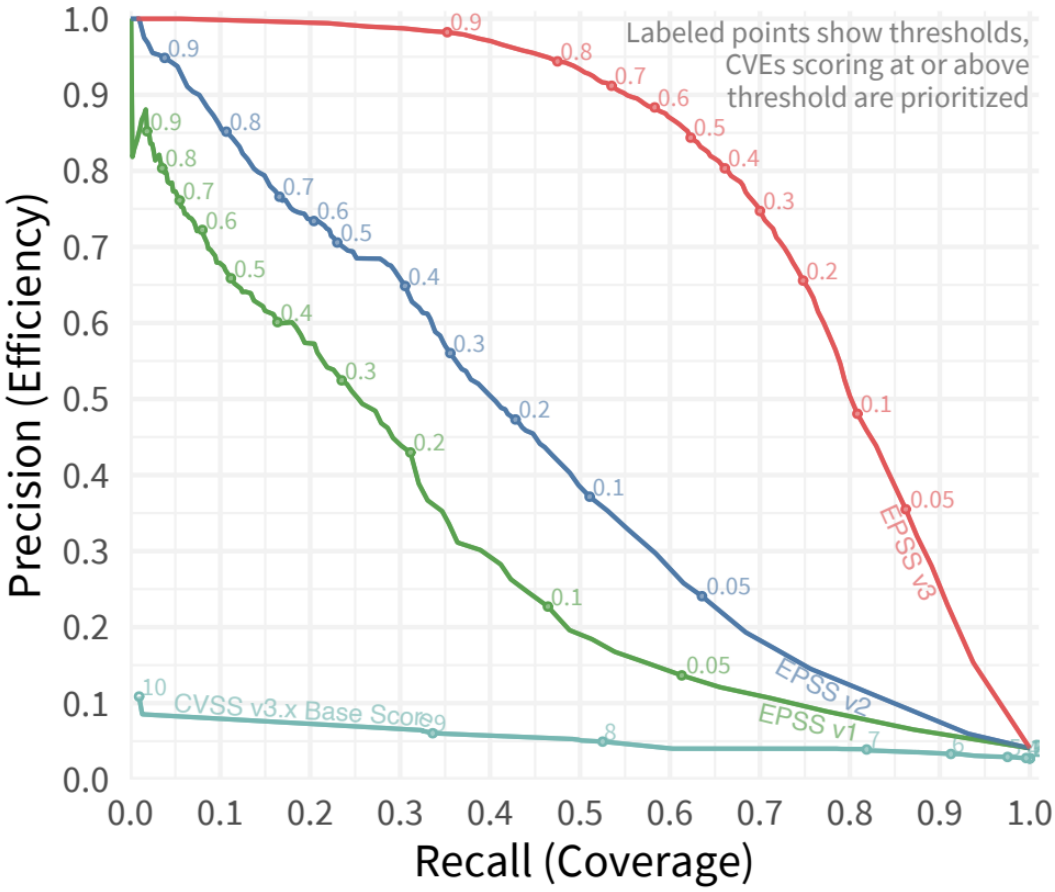

Now if we were more interested in keeping the coverage score at +-82%, what would the respective level of effort and efficiency be for either of the four approaches? This will require that the CVSS v3 score threshold is lowered to 7 and results in a level of effort calculated at 58.1% of all CVES (110,000 vulnerabilities), while the efficiency is 3.9%. EPSS v1 with a threshold of 0.015 in this scenario yields a level of effort of 44.3% accompanied with an efficiency of 7.6%. EPSS v2 with a threshold of 0.012 results in a level of effort of 39.0% of all CVEs and an efficiency of 8.9%. Repeating the exercise with EPSS v3 sees a level of effort calculated at 7.3% of all CVEs and an efficiency of 45.5% is obtained.

So, what is a desirable EPSS? According to the figures shared by the EPSS SIG an EPSS probability of 0.36 and above would achieve an efficiency of 78.5% and coverage of 67.8%. This translates into a level of effort of 3.5%, in other words 3.5% of all published vulnerabilities.

We will be investing more time in exploring the potential that seems to come from what EPSS offers. Hopefully this metric can help focus teams on being more effective in the long run.

If you want to read more details about EPSS v3 then please visit https://arxiv.org/abs/2302.14172.

[1] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid and M. Roytman, "Exploit Prediction Scoring System (EPSS)," Digital Threats: Research and Practice 2, pp. 1 - 17, 2021.
[2] J. Jacobs, S. Romanosky, O. Suciu, B. Edwards and A. Sarabi, Enhancing Vulnerability Prioritization: Data-Driven Exploit, https://arxiv.org/abs/2302.14172, 2023.
[3] J. Jacobs, S. Romanosky, I. Adjerid and W. Baker, "Improving vulnerability remediation through better exploit prediction," Journal of Cybersecurity, vol. 6, no. 1, 2020.


Ric

Snakes and Ladder Logic

You would be hard pressed to find a cyber security practitioner who wasn't familiar with the concept of vulnerability scanning. Products such as Nessus, Qualys, and even OpenVAS are practically household names nowadays. However, those scanners are all specifically targeted towards IT assets – but what about if you are responsible for the security of an operational technology (OT) environment?

Some vulnerability scanner vendors do, indeed, have special OT capability, but this generally amounts to the scanner checking typical ports used by OT protocols and being more sensitive to sending dangerous traffic that is known to cause issues in an OT environment. While those two additional facets of capability are valuable and are far safer than blindly launching a non-specialist vulnerability scanner in an OT environment, it doesn't really capture the whole picture.

In an OT environment, there is more going on than just at the network and application layers. Programmable logic controllers (PLCs), the heart of an OT environment, sense and control the physical world according to their pre-defined control logic. That control logic can be entirely bespoke, but it is more commonly built from vendor-provided 'library function blocks', and it is most often found in some form of specific language, such as ladder logic. If an adversary had access to an OT environment and wanted to cause some sort of impact in the physical world, it is most likely this control logic that they would have to target – initially for an OT-specific reconnaissance tactic called 'process comprehension', to understand what is going on in the physical world, then subsequently to manipulate values and cause the desired impact. Therefore, using a traditional port/service version-based vulnerability scanner on a PLC to discover OT vulnerabilities is a little like using one on a web server to discover web application vulnerabilities – you're looking at the wrong attack surface.

In response to this, I recently had the opportunity to work with a research team exploring what a vulnerability scanner would look like if it was built from the ground

up to target PLCs, rather than being repurposed from an existing tool which is not fit for purpose. Enter PLC-VBS, or Programmable Logic Controller Variable Block Scanner (catchy, right?), a tool which utilizes existing OT functionality to identify vulnerabilities within a PLC's control logic by scanning its memory. More specifically, PLC-VBS targets a PLC's variable blocks (areas of memory where it stores variables) and analyses the variables, and their bytes and bits therein, to determine their susceptibility to manipulation. To find more information about how PLC-VBS works technically, feel free to read the paper[5].

The work identified that of the library function blocks tested, over 90% of variables' bytes were vulnerable to being manipulated if left as their default values in the control logic – something that is common practice. This finding not only validated PLC-VBS but also highlighted the need for a focus on safe and secure PLC programming practices.

I recently wrote a blog post about the current state of OT cyber attacks[6]. While the post may have begun positively by describing a dearth of specifically OT-targeted cyber attacks, it warned that there are good conditions to foster an increase in expertise – and therefore increased frequency – of OT-targeted cyber attacks. As such, proof-of-concept techniques such as PLC-VBS should be extended and formalized to better protect OT against bespoke, targeted attacks.

Diana

Internet Crime Report 2023 – a visible age gap.

The Internet Crime Report 2022 (IC3) has just been published and shows some interesting findings. First of all, the number of complaints received during 2022 was 800,944, which is a decrease of 5% in comparison to the year before. Nevertheless, the total loss increased from being 6.9 billion USD in 2021 to 10.2 billion USD in 2022!
The number one crime type was phishing, followed by Personal Data Breach and Non-Payment/ Non-delivery (e.g. goods that were paid for but not shipped).

One very interesting finding was that of the age groups of the victims. As shown below, the age cohort suffering the highest amount of financial losses were 60+. While proportionally, the age group 20-29 had a high amount of complaints but a low amount of financial losses.

[5] https://www.sciencedirect.com/science/article/pii/S0167404823000263
[6] https://www.orangecyberdefense.com/global/blog/research/on-the-state-of-ot-cyber-attacks-and-traversing-level-35-the-artist-formerly-known-as-airgap

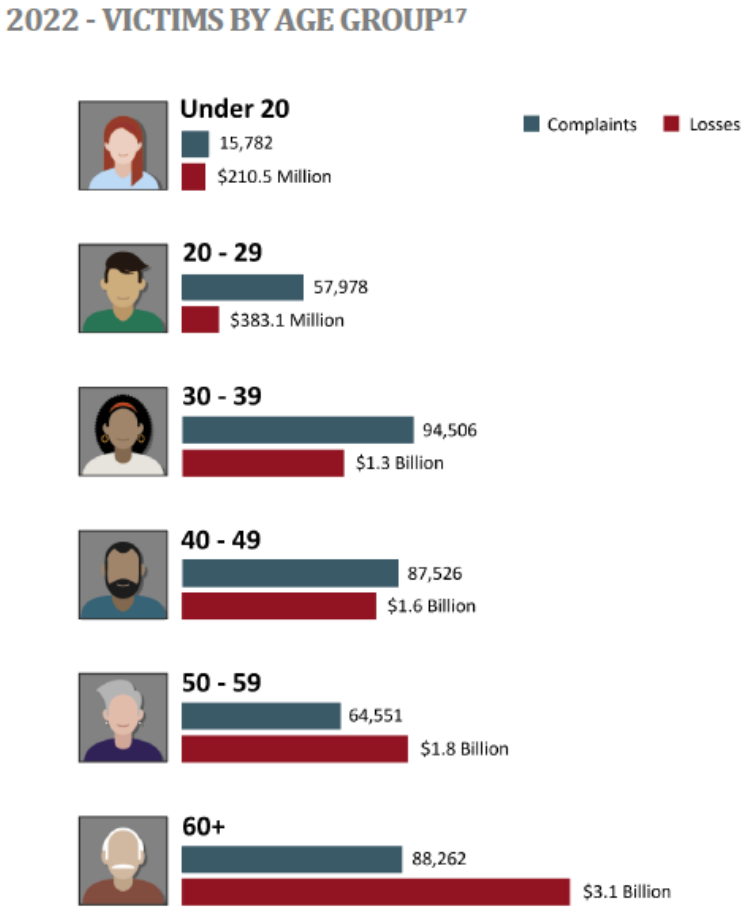

This is a very interesting finding, highlighting the gap between the 'Digital Natives' that are currently growing up with technology and all its advances and disadvantages and the 'Silver Surfers' as they are sometimes referred to, that have adapted and learned how to use technology but seem to fall victim with very high financial losses. A few years ago, I conducted a study that looked at exactly these two 'generations' and forms of interpersonal cybercrime. The findings were very similar, the younger generation fell similarly victim to these crimes but did not show such a high financial loss. One explanation is that the elderly often lose a great amount of their savings to fraudulent activities, while the younger generation does not necessarily have this kind of capital (yet).

A second potential explanation is *trust*. The elder generation generally trusts human interaction much more (it's something they know well from the offline-life) and thus they become more vulnerable to these schemes. Digital natives have most likely become victim of some form of anti-social / harmful behavior online in their first years of technology usage, and have partially learned through bad experience not to trust individuals online.

And while they don't have the life experience and wisdom of the elder generations, their (online) behavior tends to be more careless, and thus might cause victimization. But their understanding of technology is much better, and thus, the impact (in this case financial loss) might be lower.

And lastly, while the report collects most of the complaints form the U.S., cybercrime is a global problem and impacts us globally. The top 10 victim countries outside of the U.S. are the United Kingdom, Canada, India, Australia, France, South Africa, Germany, Brazil, Mexico and the Philippines.
The full report can be found here: https://www.ic3.gov/Home/AnnualReports

Good News Cyber

Dutch police arrest three ransomware actors

Dutch police have arrested three men allegedly behind ransomware attacks which impacted thousands of companies, the attacks culminating in ransom demands of between €100,000 and €700,000, with one of the suspects believed to have made over €2.5 million.

It is thought that tens of millions of pieces of personal information were stolen by the trio, including names, email addresses, telephone numbers, bank account numbers, credit card details, account passwords, license plates, and passport details. The Dutch police stated that even when organisations paid the demanded ransom the stolen data was still sold on dark web marketplaces to generate extra profit.

It has also been reported that one of the trio arrested was an ethical hacker who was a member of the Dutch Institute for Vulnerability Disclosure (DIVD), a group established by the Dutch government to hunt for vulnerabilities in computer systems. Media reports claim that this researcher abused the access they had to vulnerability details in order to help facilitate the ransomware attacks.

Collaboration between German police & Ukrainian police & Europol & FBI targeted core members from DoppelPaymer

In co-ordinated actions the German Regional Police and the Ukrainian National Police, with the assistance of Europol, the Dutch police and the United States FBI, arrested suspected core members of the group believed to be behind the DoppelPaymer ransomware.

During the arrests German officers raided the house of a German national, who is believed to be a senior player in the DoppelPaymer ransomware group. At the same time, despite the difficulties posed from the Russian invasion, Ukrainian officers arrested a Ukrainian national who is also believed to be a member of the core DoppelPaymer group. In both cases electronic equipment was seized for forensic analysis to help determine the exact roles each suspect had within the ransomware group.

Australian led International Counter Ransomware Task Force commenced operations

In January Australia began the operations of the International Counter Ransomware Task Force. The task force forms part of the US-led Counter Ransomware Initiative which is made up of 37 different governments from around the globe.

The aim of the task force is to disrupt, combat and defend against the ever present ransomware threat. This goal will be achieved through international cooperation consisting of information and intelligence exchanges, sharing best practice policy and legal authority frameworks, and collaboration between law enforcement and cyber authorities.

The current members of the Counter Ransomware Initiative are: Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Croatia, Czech Republic, Dominican Republic, Estonia, France, Germany, India, Ireland, Israel, Italy, Japan, Kenya, Lithuania, Mexico, the Netherlands, New Zealand, Nigeria, Norway, Poland, Republic of Korea, Romania, Singapore, South Africa, Spain, Sweden, Switzerland, United Arab Emirates, United Kingdom, United States, Ukraine and the EU.