Orange
Cyberdefense

Security Intelligence | Research Report

# Monthly Report
# May 23

# Contents

# Introduction

### RDP Honeypot Targeted 3.5 Million Times in Brute-Force Attacks

Researchers at GoSecure, a threat hunting and response company, conducted an experiment using high-interaction honeypots with RDP connections accessible from the public web. They recorded nearly 3.5 million login attempts to their RDP honeypot system over three months. The honeypot has been functioning for over three years and running steadily for over a year. The attack count for the entire year reached 13 million login attempts. Researchers named the system to appear like a bank's network, and the compromise attempts relied on brute-force attacks. In some 60,000 cases, the attacker did reconnaissance before trying the right login.

### Fortinet Fixes Critical Pre-Authentication Remote Code Execution Vulnerability in SSL-VPN Devices

A highly critical vulnerability responsibly reported by 2 researchers and disclosed by the vendor in an Advance Warning to their partners is now patched by Fortinet.

Tracked as CVE-2023-27997, the flaw enables remote code execution (RCE) but does not require the attacker to be logged in. This vulnerability is most probably present in all Fortinet SSL VPN appliances, as this bug affects all previous versions. According to Shodan, over 250,000 Fortigate firewalls are accessible from the Internet.

### A Zero-Day Vulnerability Affecting MOVEit Exploited in Ongoing Attacks

A zero-day vulnerability affecting popular file transfer tool MOVEit sold by Progress Software has been leveraged by hackers to steal data. All versions of the product are impacted by this critical flaw, so it is highly recommended to apply the patch released by the company on May 31st in their advisory as soon as possible, and hunt for signs of compromise if you use the solution.

Microsoft officially attributed the ongoing exploitation of this critical flaw to Lace Tempest (a.k.a. DEV-0950 according to their previous nomenclature), a known ransomware affiliate of the Cl0p Cy-X group.

Horizon3 has publicly shared an exploit for the vulnerability, tracked as CVE-2023-34362, as well as a list of malicious patterns in the database tables that could be hunted for. As this exploit can work in default MOVEit Transfer configurations, it is likely that more threat actors may attempt to carry out attacks against unpatched servers left exposed online.

**At a glance**

Our research teams have been tracking Cyber Extortion (Cy-X) activity for the last few years and have collected information on over 6500 victims to date. Our 'Cy-Xplorer' report summarizes the findings regarding these victims, the criminals and the tools and techniques they use.

Download our Cy-Xplorer report HERE.

# World Watch Review

The Orange Cyberdefense CERT published a total of 15 new World Watch advisories during May 2023, along with adding updates to a further 19 previously published advisories.

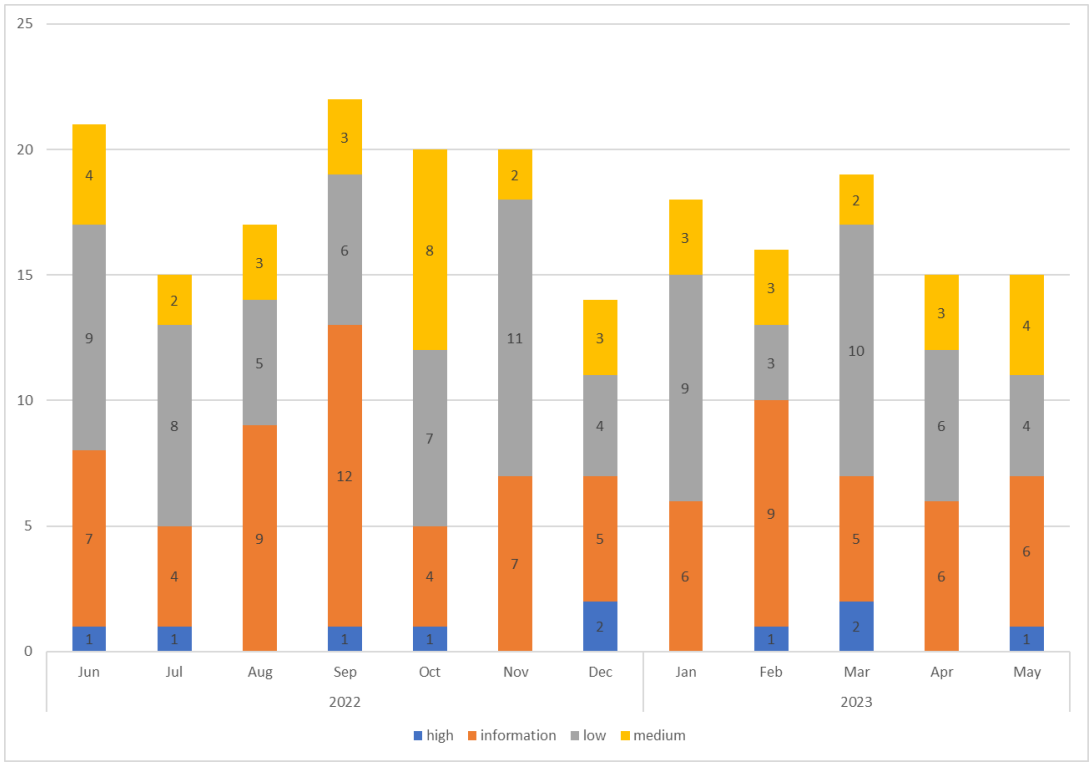


Breakdown of Published Advisories Previous 12 Months

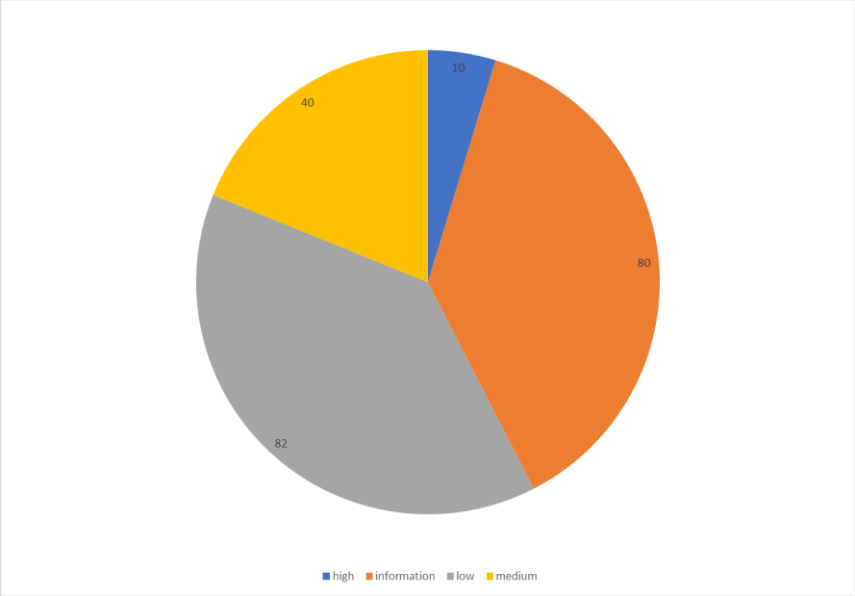


Breakdown of Advisory Criticality for Previous 12 Months

## Advisory Summary

As can be seen above there was just 1 advisory rated as high criticality during May with all other 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.

See below for a timeline of advisories rated Medium and higher:

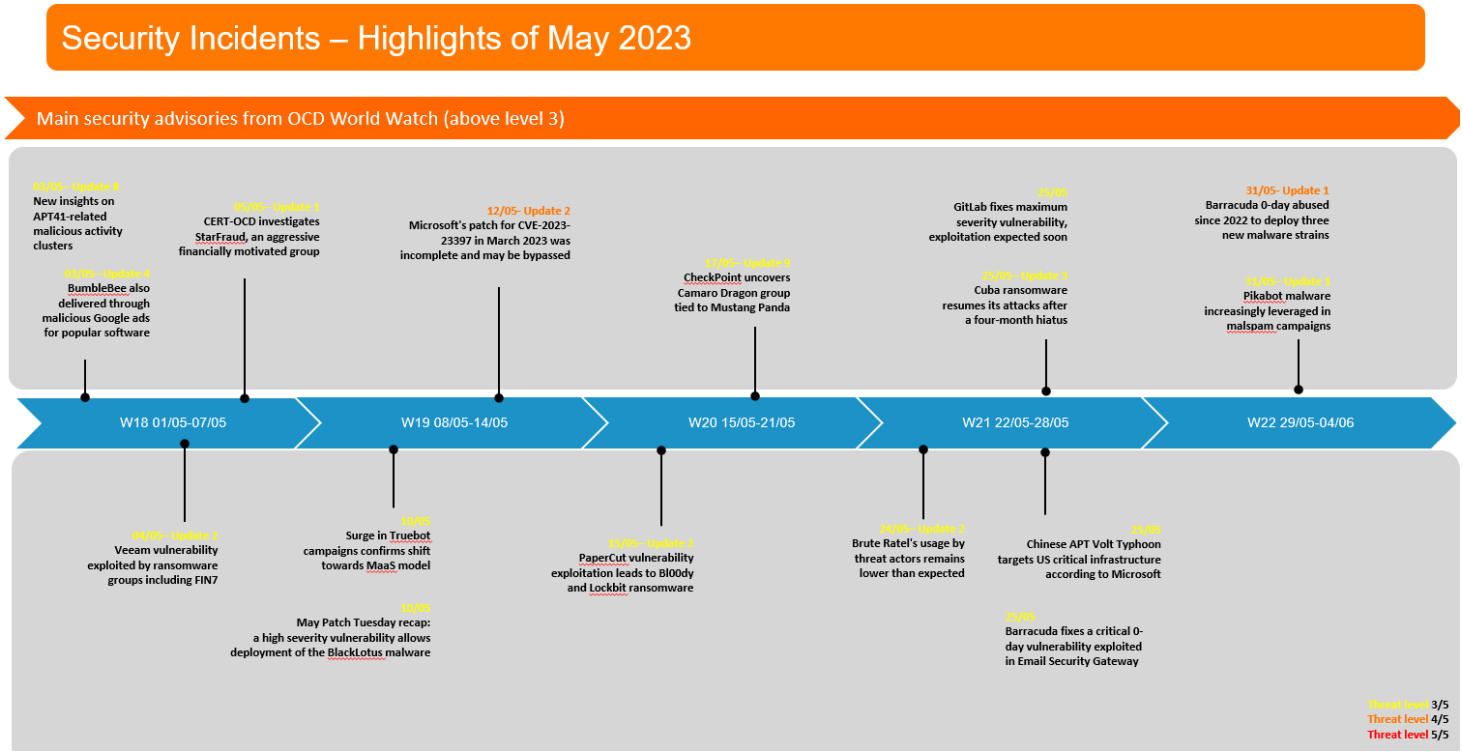

# Editor's Notes

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.

Wicus

## The Secure Artificial Intelligence Framework

In early June 2023 Google published a conceptual framework called Secure Artificial Intelligence Framework (SAIF). As the name suggests, it seeks to put forward best practices that guide practitioners down a path to use Artificial Intelligence (AI) in a manner that yields outcomes generally accepted by society that is in line with societal norms. SAIF is not unique in its concept or intent but builds or leverages existing schools of thought on dealing with AI in a way that is risk sensitive.

NIST published the first version of the AI Risk Management Framework (AIRMF) in January 2023 and seeks to provide a means to define a risk management approach based on risk tolerance and prioritization that can be used to establish trustworthy AI characteristics. The characteristics include discrete facets such as being Safe, Secure and Resilient, Explainable and Interpretable, Privacy Enhanced, and Fair-With Harmful Bias Managed. These are kept in check by ensuring the facets are Valid and Reliable by requiring the AI system to be Accountable and Transparent.

Where the NIST framework seeks to be more risk centric, Google's SAIF leans more toward the pragmatic and aims to be a guide on how to start thinking about the use of AI technology in a business in terms of established security practices. The former is much more formal and rigorous while the latter follows a less formal tone and is more actionable initially. To put it differently, SAIF is more accessible to a wider audience, where AIRMF requires a bit more intellectual investment to get to grips with. Not only does SAIF describe much of what is outlined in AIRMF at a higher level, but SAIF bridges the gap between theory and practice. The latter is illustrated by a four-step plan.

SAIF's four-step plan starts with understanding the use case of AI, hence the name "Understand and use". As the name suggests, it's about knowing if and why AI solutions are needed as well as what type of AI solutions should be considered for a specific task. Also, will AI models be exposed externally and work with untrusted input.

The second step involves establishing a cross functional team for the specific AI use case and is named "Assemble the team". This team can consist of stakeholders from a wide range of disciplines and include the business use case owners, security teams, cloud engineering teams, risk and audit professionals, privacy experts, legal experts, data scientists, developers, and responsible AI and ethics experts.

The third step is named “Level set with an AI primer” and seeks to ensure that all non-AI experts have a good understanding of the basics of how AI models are constructed, functions, and the limitations of the AI solution. This step does not seek to make experts of all the team members but emphasizes the need to ensure that everyone involved understands at a high level the basic workings of the AI model.

The final step involves applying the six core elements of SAIF namely:

1. Expand strong security foundations to the AI ecosystem.
2. Extend detection and response to bring AI into an organization’s threat universe.
3. Automate defenses to keep pace with existing and new threats.
4. Harmonize platform level controls to ensure consistent security across the organization.
5. Adapt controls to adjust mitigations and create faster feedback loops for AI deployments.
6. Contextualize AI system risk in surrounding business processes.

Each of these six elements contains several subpoints and I touch on some of them.

On point one, Expand strong security foundations to the AI ecosystem, supply chain management is brought into the equation here, specifically with regard to the supply chain assets involved, training data, and code that was used to create the AI model. This touches on a level of expertise that knows data governance and spans data quality, data security, data architecture, metadata, data lifecycle management, and data storage. Point one also states that the business must try to retain and continuously train staff to ensure a strong and secure AI ecosystem is cultivated.

Point two, Extend detection and response to bring AI into an organization’s threat universe, is a challenge and opportunity for security researchers to find new ways to monitor AI systems for potential exploitation. This will require those that monitor these systems to understand the capabilities of these models and what kinds of abuse or exploitation is possible. This will also require staying abreast of the latest techniques used to manipulate these kinds of models.

Point three, Automate defense to keep pace with existing and new threats, can span the traditional security mindset and flow into new adversarial approaches that introduce force multipliers for automation and efficiency. At the same time this could lead to more effective defenses that can be used to counter attackers.

Point four, Harmonize platform level controls to ensure consistent security across the organization, involves the tooling and use of frameworks dealing with AI in the

business. Here businesses should seek to prevent further fragmentation that typically comes with adoption of new technologies.

Point five, Adapt controls to adjust mitigations and create faster feedback loops for AI development, shows that existing security teams can use what they already know to target systems using AI. One such example is Red Teaming exercises to verify the safety and security of such systems. This includes identifying new novel attacks, such as prompt injection, or using the latest state of the art tactics or techniques to probe for AI system vulnerabilities. Another interesting aspect here is that SAIF recommends creating a feedback loop for this element to improve the red team's future capabilities but also to ensure that what is learned can be incorporated in other training data sets to adjust future iterations of the model.

Finally, point six, Contextualize AI system risk in surrounding business processes, speaks to ensuring that the business has people that understands AI-related risks and how to deal with them. This element also speaks to the management of AI model inventories with risk profiles based on specific use cases when using third-party solutions or services. Similar to the NIST AIRMF the idea is to establish a risk tolerance for AI use cases, meaning that the business knows the associated risks and knows how to deal with them to limit impact.

This was a condensed overview of the SAIF and it is best to explore the SAIF documentation in full. Consider spending the extra time also reading through NIST AIRMF documentation. SAIF is not a complete primer for AIRMF, but it will prepare you for some ideas that are present in the NIST documentation. AIRMF is more grounded in formal methodologies and tied into existing standards and may be suitable for enterprise businesses as it can tie in formally with existing governance and risk management practices. Google's SAIF does feel more attainable for the average business. It is likely that some government regulation will be created to ensure responsible and accountable use of AI solutions in the future, much like what was done for how businesses can use personal data.


Ric

## 6 Month Roundup

This is my 6th month in the Orange Cyberdefense Security Research Center and so I thought I'd mark the occasion with a small, retroactive introduction, along with a recap of the research I have been publishing and working on so far. I will provide links to the work, where possible. However, if you think any of the work is of particular interest and you'd like to discuss it further please do find me on social media or contact the research team to reach me via email.

I have been working in cyber security for around 10 years in a consultancy capacity. However, I began more serious research work during my PhD. The core focus of my work during that time was quantitative cyber risk assessment (my

thesis in particular[1]). I was introduced to operational technology (OT) while at the university and have been enthralled by it ever since. After the PhD I had a role working in OT cyber security, where I had the opportunity to see the technology firsthand; but the call of research was too strong, with too many questions unanswered, and too many ideas unimplemented. Therefore, in December 2022 I joined the Security Research Center at Orange Cyberdefense!

My first output in the role came in January, whereby work I contributed to was accepted to the journal Computers and Security[2]. Regular readers will have seen this work discussed in my editorial in March. This work implemented a proof-of-concept vulnerability scanner that can identify vulnerabilities within a programmable logic controller's (PLC's) control logic. Rather than the typical network-level vulnerability scanners used today, that are reappropriated from IT environments, this work looked deeper into the programming practices and memory management of PLCs to identify how they could contribute to an attack on an OT environment.

The previous work was carried over from earlier research and it was simply accepted while I was in this role. The first piece of work I started while at Orange Cyberdefense came in February, in the form of a blog post[3]. With this post, I aimed to dispel the myth that OT is constantly the target of attacks by looking at the state of historical and contemporary attacks, as well as describing the challenges adversaries encounter when conducting attacks that specifically target OT. However, the positive message was short lived as I went on to describe how there is a storm afoot, as adversaries are showing an appetite for attacking OT and the contextual knowledge is becoming more accessible for them to quickly develop their capabilities.

From March until the time of writing I have been writing papers for academic conferences. The first paper to have been accepted was presented by a co-author on Monday the 12th June in Barcelona, Spain, at the IFIP Networking 2023 workshop "Impact of IT/OT Convergence on the Resilience of Critical Infrastructures"[4]. This paper (accessible 2023-06-15[5]) looks at the current state of OT standards and guidelines and benchmarks them against existing, mature IT standards and guidelines. I wrote about this paper in my previous editorial in May. The second paper to have been accepted will be presented by myself on Friday 7th July in Delft, Netherlands, at the EuroS&P 2023 workshop "Re-design Industrial Control Systems with Security"[6]. This work addressed the fact that, at present, most PLCs are programmed using library functions supplied by vendors, which cannot be viewed or edited and therefore cannot have secure PLC programming

[1] https://eprints.lancs.ac.uk/id/eprint/172697/2/2022derbyshirephd.pdf
[2] https://www.sciencedirect.com/science/article/pii/S0167404823000263
[3] https://www.orangecyberdefense.com/global/blog/research/on-the-state-of-ot-cyber-attacks-and-traversing-level-35-the-artist-formerly-known-as-airgap
[4] https://networking.ifip.org/2023/index.php/program
[5] https://eprints.lancs.ac.uk/id/eprint/192592/
[6] https://ricssworkshop.github.io/program.html

practices applied. The paper took the stance that open-source functions may be the best alternative, as they can be viewed and edited. However, after applying the vulnerability scanner from our previous work, we confirmed that these open-source functions were still vulnerable. Therefore, a framework was put forward for a community initiative to ingest and secure open-source PLC code. Unfortunately, as the conference has not taken place, a paper is not currently available. If you're interested, please do get in touch for a preprint. The final piece of work is still under double blind review, and as such it wouldn't be prudent to discuss it. However, once the work is published (at the current venue or another, should it be rejected – which is fine), I will be writing about it in much detail in a blog post and hopefully giving a few talks, too.

During the time I spent writing papers, I have also been involved in a number of speaking engagements. In March I presented at BSides Lancashire about the current state of OT cyber attacks, which covered much of the content from my January blog post[7]. In April, I contributed to a panel online for the Think Cybersecurity for Government event, discussing cyber security within the supply chain and all the challenges involved[8]. In May, I gave an online presentation to the Cyber Wales OT Cluster about how to exploit PLC programming practices and memory management for enumeration, exploitation, and command & control[9]. Finally, this month I will be speaking on a panel at the PETRAS Living Securely in the Internet of Things Event. The panel will be focusing on holistic cyber security for operational technology in critical national infrastructure[10].

Next on the agenda, other than contributing to this year's Security Navigator, I have a few more OT-focused projects I'd like to tackle, time permitting; one which is particularly technical and two which take a look at the bigger picture of OT cyber security culture for both those working in the field to defend it and the adversaries.


Diana

## DDoS & Ransom (& hacktivism)

May has been an interesting month when it comes to cyber activism/hacktivism! Especially the Nordics, Sweden & Denmark were impacted by yet another threat actor group called NoName057(16). In our last (April) monthly report, I took the time to explain what we are currently observing, focusing mainly on Anonymous Sudan and NoName057(16).

During May, Anonymous Sudan began to demand ransom from the Scandinavian airline SAS[11]. Anonymous Sudan has increased their initial demand from $3500

[7] https://www.youtube.com/watch?v=H1TzFVhdoag
[8] https://www.thinkdigitalpartners.com/news/2023/05/17/tackling-cybersecurity-in-the-supply-chain/
[9] https://cyberwales.net/events/?event=ot-cluster-may23
[10] https://petras-iot.org/update/petras-iet-event-2023-living-securely-in-the-internet-of-things/
[11] https://therecord.media/hacker-group-anonymous-sudan-demands-three-million-from-sas

USD to over $10 million. How this amount can be appropriate (if it ever could be) stands to question. While ransoming via DDoS (RDoS) is nothing new, it is interesting to observe a former hacktivist group that had political/ideological motives to attack anyone who opposed Sudan/Islam; to now turn financially motivated, and thus we cannot consider them as hacktivists any longer.
This is of course one way to do it. Another incident was experienced and reported on by Imperva[12] in the beginning of March 2022. Here the group that masqueraded as REvil dropped several ransom notes BEFORE starting to DDoS the victim, demanding 1 BTC per day.

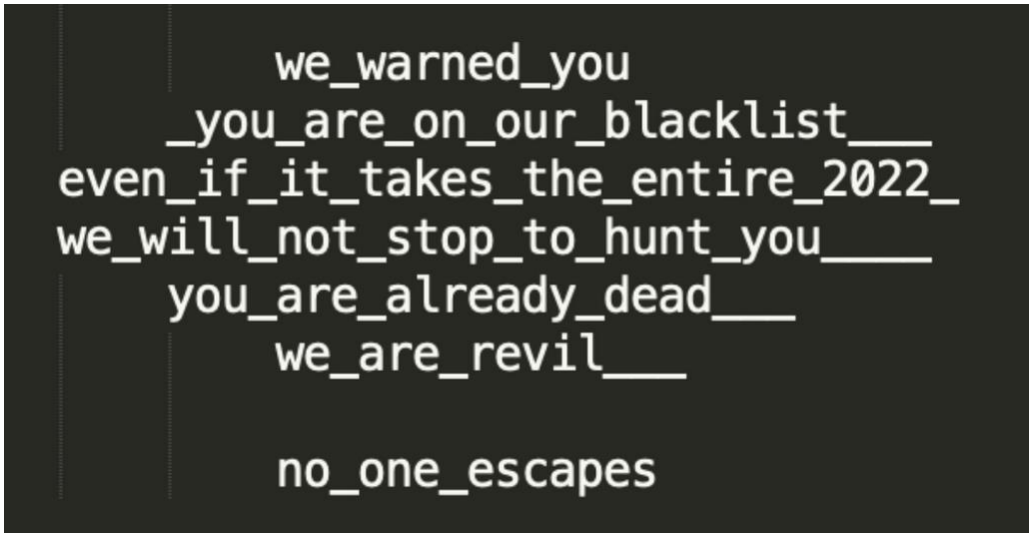


This modus operandi has a very low entry level, given the fact that no compromise is required; especially when DDoS is available as a service component like other forms of cybercrime.

[12] https://www.imperva.com/blog/imperva-mitigates-ransom-ddos-attack-measuring-2-5-million-requests-per-second/

This week we received updated hacktivist data via our partnership with Intel471. Sweden remains on the second top most impacted countries in 2023.

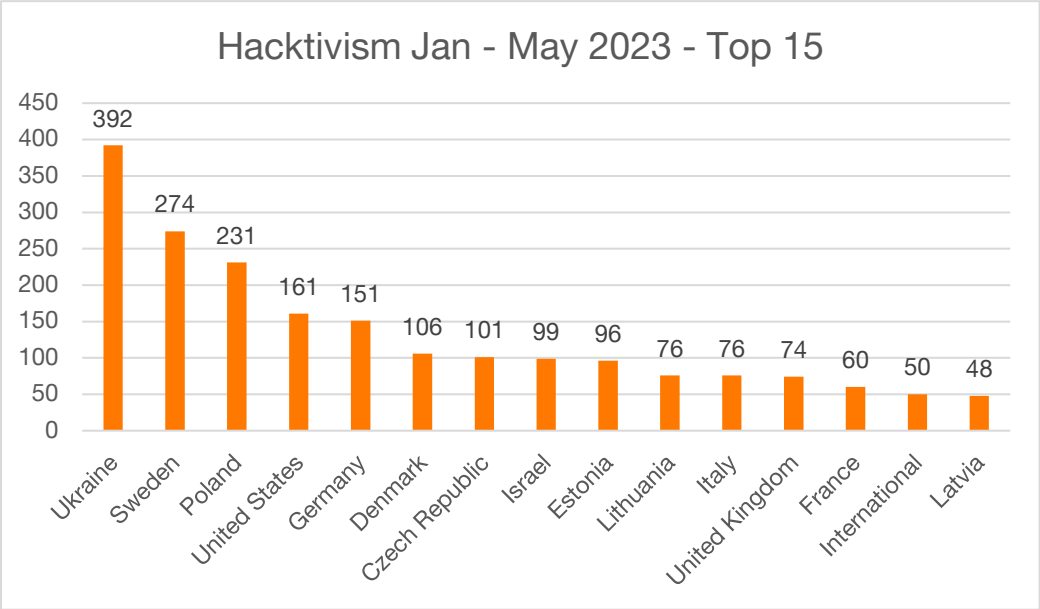


With a few volunteers from within OCD, we are currently developing our own dataset, where we will start collecting and structuring data based on the Telegram channel announcements relevant hacktivist groups are making. Here we are planning to collect date stamp, country, industry & hacktivist groups.

# Good News Cyber

- Microsoft says SMB signing (aka security signatures) will be required by default for all connections to defend against NTLM relay attacks. This security mechanism has been available for a while now, starting with Windows 98 and 2000, however due to slow SMB data transfers it was never enabled by default, it has though been updated in Windows 11 and Windows Server 2022 to improve performance and protection by significantly accelerating data encryption. Microsoft Principal Program Manager Ned Pyle said, "Expect this default change for signing to come to Pro, Education, and other Windows editions over the next few months, as well as to Windows Server."

- Apple recently unveiled its newest privacy and security advancements, which include significant changes to Lockdown Mode, Communication Safety, and Private Browsing in Safari. Additionally, Apple introduced brand-new features including Check In, NameDrop, and Live Voicemail, all of which were created with privacy and security in mind.

- Following a number of events over the past few weeks, the Python Software Foundation has taken various steps to enhance the security and privacy of the official Python Package Index (PyPI). They are working on reducing the requirement for the PyPI portal to store a user's IP address and are enabling two-factor authentication (2FA) for PyPI accounts. By the end of the year, all accounts that maintain Python libraries on the PyPI site will have their access to specific PyPI features restricted if they do not configure a 2FA method.

- A Manhattan federal court gave Romanian national Mihai Ionut Paunescu, aka "Virus," a three-year prison term for running a bulletproof hosting service and assisting the spread of malware.

  Web hosts that offer "bulletproof hosting services" have permissive regulations regarding their customers' illegal content and behaviour and are often located in nations with lax or non-existent internet laws. These services are also notorious for disobeying takedown demands from authorities and owners of copyrights.

  The Department of Justice claims Paunescu's service made it easier for DDoS (distributed denial of service) attacks to be launched, spam to be sent around the world, and a number of information-stealing and banking malware families, including Gozi (Ursnif), Zeus, SpyEye, and BlackEnergy to spread.

- The number of phishing websites tied to domain name registrar Freenom plummeted in the months following a recent lawsuit from social networking giant Meta, which accused Freenom of ignoring abuse complaints about phishing websites while monetizing traffic to them. Freenom has never charged for domains in the five country code top level domains they manage, but the registrar maintains the right to take them back and redirect traffic to other sites. Freenom was responsible for over half of all new phishing domains from country-code top-level domains by December 2022, when Meta filed its case. Meta's bid to seal its Freenom case was denied. Meta resubmitted their December 2022 complaint in March 2023.

“Freenom provides free domain name registration services and shields its customers’ identity, even after being presented with evidence that the domain names are being used for illegal purposes,” Meta’s complaint stated.

Meta cited Interisle Consulting Group's 2021 and 2022 findings indicating Freenom's five ccTLDs made up half of the Top Ten Phishing TLDs.

Interisle partner Dave Piscitello said something remarkable had transpired in the months since the Meta case. “We’ve observed a significant decline in phishing domains reported in the Freenom commercialised ccTLDs in months surrounding the lawsuit,” Piscitello remarked on Mastodon. “Responsible for over 60% of phishing domains reported in November 2022, Freenom’s percentage has dropped to under 15%.”