WITH secure

Threat Highlight Report

January 2025

Table of Contents

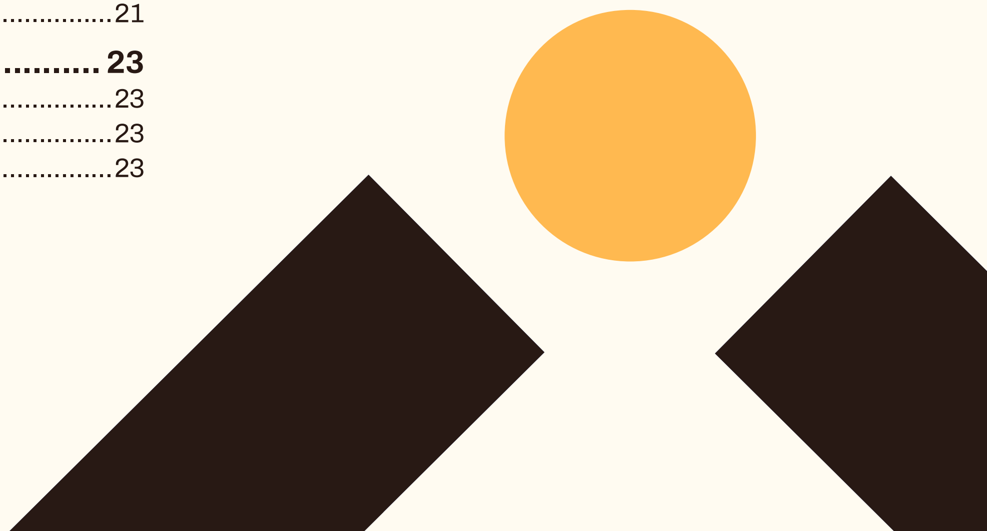

Foreword

“There was a distinct theme of incidents and stories around Identity this month, which is possibly to be expected when a recent report finds that SaaS breaches have increased 300% year on year. There are details of a campaign targeting PayPal accounts via abuse of legitimate Microsoft services and PayPal’s very unintuitive account management methodology, and details of a flaw in Google’s OAuth implementation which enables the resurrection of defunct domains and email addresses, and the access of SaaS accounts that were associated with those email addresses.

It would be difficult to write about January 2025 without mentioning DeepSeek, so we haven’t tried to.

There has also been continued news regarding mass exploitation and vulnerable edge devices, with a number of concerning findings coming from an investigation by Eclypsium into the hardware and software of Palo Alto firewalls. While their investigation was only into one brand of firewall, considering how many security incidents there have been due to edge network security devices then similar issues could probably be found in most such devices.

The THR this month is rather more brief, so as to make it easier to consume and simpler for readers to find stories relevant to their interests.

As ever we will be discussing the content of the THR in the Cyber Threats Xposed podcast, so if a podcast is an easier format for you to access and digest, please do find us on your podcast platform of choice.”

Stephen Robinson,
Senior Threat Intelligence Analyst, Threat Intelligence and Outreach, WithSecure

Monthly highlights

Identity

Research from Obsidian suggests that SaaS breaches increased 300% in 2024 compared to 2023. While this may sound drastic, it is, all things considered, a reasonable number. Use of SaaS has continued to grow rapidly, and so has targeting by attackers.

Sometimes attackers take advantage of unexpected quirks in the complex cloud ecosystems, a recent interesting example of this was identified by researchers at Trufflesecurity, who found that a flaw in Google's OAuth "Sign in with Google" feature allowed anyone to register lapsed or defunct domain names, recreate previously existing email addresses, and access any SaaS accounts associated with them through the Sign on with Google feature. This is possible because the Google OAuth sub-claim system has a 0.04% "inconsistency", i.e. failure rate. While this is relatively low, when millions or billions of authentications are being made, it adds up rather quickly. As such, downstream consumers of Google OAuth verification cannot rely upon the sub-claim, and instead have to rely solely on the email address and the hosted domain.

Social engineering to gain access to SaaS accounts is also a big issue, particularly when attackers can create their infrastructure on a provider's cloud service, then use that "legitimate" infrastructure to target other customers. A campaign that was documented by Fortinet Labs this month involved the attack technique of registering a Microsoft test domain, then creating a distribution list on that domain and adding victim email addresses to it. An attacker can then make a PayPal payment request directed at that distribution list, and the PayPal service will automatically send a legitimate email to the distribution list, which will be forwarded to the victim. Because the email comes from a legitimate Microsoft domain, it will bypass most standard security checks. Nothing actually happens however, until the victim clicks on the link within the email and views the payment request. At that point, PayPal will automatically associate the distribution list to which the payment request was sent with the victim's account, allowing the attacker who controls the distro list to login to the victim account and control it as they wish. All of which occurs without any notification to the user, it is simply an automatic action on the part of PayPal, which seems to have simply not considered that emails can be forwarded from one email address to another.

Monthly highlights

Identity

Research from Obsidian suggests that SaaS breaches increased 300% in 2024 compared to 2023. While this may sound drastic, it is, all things considered, a reasonable number. Use of SaaS has continued to grow rapidly, and so has targeting by attackers.

Sometimes attackers take advantage of unexpected quirks in the complex cloud ecosystems, a recent interesting example of this was identified by researchers at Trufflesecurity, who found that a flaw in Google's OAuth "Sign in with Google" feature allowed anyone to register lapsed or defunct domain names, recreate previously existing email addresses, and access any SaaS accounts associated with them through the Sign on with Google feature. This is possible because the Google OAuth sub-claim system has a 0.04% "inconsistency", i.e. failure rate. While this is relatively low, when millions or billions of authentications are being made, it adds up rather quickly. As such, downstream consumers of Google OAuth verification cannot rely upon the sub-claim, and instead have to rely solely on the email address and the hosted domain.

Social engineering to gain access to SaaS accounts is also a big issue, particularly when attackers can create their infrastructure on a provider's cloud service, then use that "legitimate" infrastructure to target other customers. A campaign that was documented by Fortinet Labs this month involved the attack technique of registering a Microsoft test domain, then creating a distribution list on that domain and adding victim email addresses to it. An attacker can then make a PayPal payment request directed at that distribution list, and the PayPal service will automatically send a legitimate email to the distribution list, which will be forwarded to the victim. Because the email comes from a legitimate Microsoft domain, it will bypass most standard security checks. Nothing actually happens however, until the victim clicks on the link within the email and views the payment request. At that point, PayPal will automatically associate the distribution list to which the payment request was sent with the victim's account, allowing the attacker who controls the distro list to login to the victim account and control it as they wish. All of which occurs without any notification to the user, it is simply an automatic action on the part of PayPal, which seems to have simply not considered that emails can be forwarded from one email address to another.

WithSecure Insight

More and more malicious actors are not just cloud-able but cloud-focused, because legitimate enterprises are becoming more and more heavily invested in the cloud. While the hybrid cloud/on premise technologies of the modern enterprise can be affected by on-premise compromises and activity, when targeting the cloud actors are able to take advantage of some of the same economies of scale that legitimate users are benefitting from. In addition to this, cloud environments rely upon identity, something which attackers have shown again and again they can easily steal if victims do not have adequate protections in place. Once an attacker has the necessary identity credentials, they are then free to monetize a cloud environment however they wish. Finally, cloud environments are complex, and in modern SaaS deployments there can be many different layers and relations of suppliers and customers, each with their own configurations and quirks.

DeepSeek

The launch of the Chinese LLM DeepSeek has shaken up the LLM industry, wiping huge sums of money off the share price of some companies. However, to us that is less interesting than the security implications. In the first week that it was launched, DeepSeek announced that they were experiencing disruption due to large scale malicious attacks, which most likely means DDoS. They may have been experiencing a DDoS, but as yet there is no indication that it was a malicious attack, it may simply have been caused by overwhelming demand for their service. Soon after, researchers at Wiz discovered that DeepSeek had left a database containing millions of lines of internal data and sensitive information completely publicly exposed, allowing full control to anyone who found it. The exposed data include chat histories, information about the backend data and services, operational details, and even API secrets. Researchers at Wiz believe that the level of exposure was such that, even without leveraging the exposed secrets within the database, it would have been possible to escalate privileges and perform malicious activities within the DeepSeek environment.

WithSecure Insight

More and more malicious actors are not just cloud-able but cloud-focused, because legitimate enterprises are becoming more and more heavily invested in the cloud. While the hybrid cloud/on premise technologies of the modern enterprise can be affected by on-premise compromises and activity, when targeting the cloud actors are able to take advantage of some of the same economies of scale that legitimate users are benefitting from. In addition to this, cloud environments rely upon identity, something which attackers have shown again and again they can easily steal if victims do not have adequate protections in place. Once an attacker has the necessary identity credentials, they are then free to monetize a cloud environment however they wish. Finally, cloud environments are complex, and in modern SaaS deployments there can be many different layers and relations of suppliers and customers, each with their own configurations and quirks.

Mass Exploitation

In a short but significant story month, Eclypsium have performed a hardware and software assessment of Palo Alto Next Generation Firewalls (NGFWs) and found what they describe as significant issues. While these are intended to be highly secure, reliable, enterprise devices, the researchers found that they were built with commodity hardware running known vulnerable software and firmware, and they were missing certain expected, foundational security features. This included configuration flaws within Palo Alto's Secure Boot implementation which would allow an attacker to simply bypass Secure Boot entirely, allowing for the installation of an undetectable boot kit.

WithSecure Insight

Attackers would need to gain administrative access to these devices in order to exploit some of the vulnerabilities identified, however with an ongoing stream of RCE and privilege escalation vulnerabilities being discovered in network infrastructure, this is not a particularly tall order, especially considering that network appliances are often not rapidly patched. Really though, the concern this research raises is not the specific vulnerabilities themselves, but the fact that multiple known vulnerable pieces of software are included in what is supposed to be a black-box security device.

Ransomware

The following data is limited to multi-point of extortion groups who are operating a leak site which is parseable. In this section we will be looking at victim leak sites. This dataset is probably the best and most consistent source we have that enables us to understand the landscape, but the data collected here is not fallible, there are several variables that impact and skew this dataset:

- It is attacker led, and some attackers may be incentivized to post incorrect data.
- It is fluid, and victims are added and removed frequently.
- Extortion success is another key factor, if the amount of paying victims greatly increases, 'total' ransomware numbers may also appear to decrease.

With this said, we can draw some insight from this data with sensible assumptions – and recognizing the data isn't perfect, it does provide a decent gauge on the ransomware landscape. The assumptions the industry typically abide by are:

- There is a roughly relatively consistent month-on-month victim payment rate,
- Actor posts do contain an element of truth.

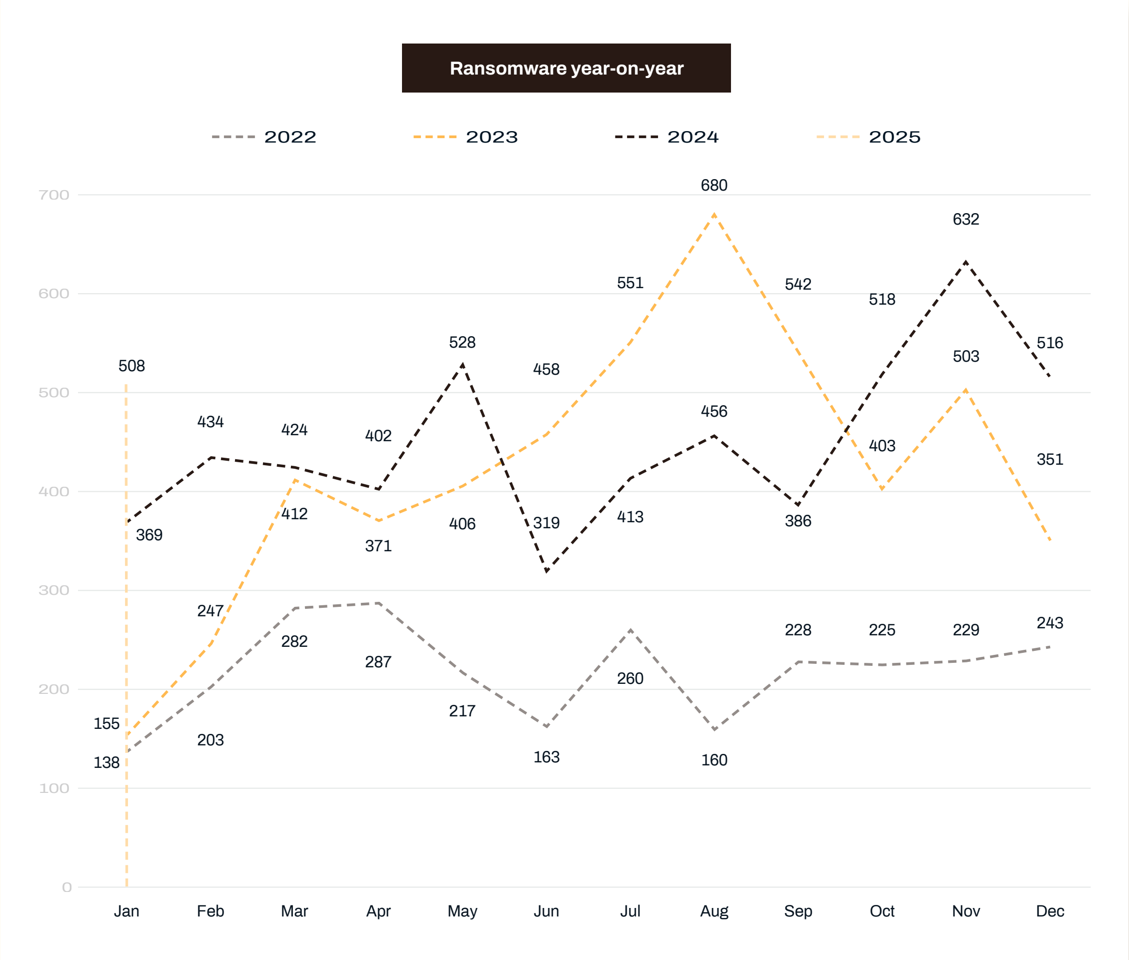
Ransomware year-on-year

Ransomware

The following data is limited to multi-point of extortion groups who are operating a leak site which is parseable. In this section we will be looking at victim leak sites. This dataset is probably the best and most consistent source we have that enables us to understand the landscape, but the data collected here is not fallible, there are several variables that impact and skew this dataset:

- It is attacker led, and some attackers may be incentivized to post incorrect data.
- It is fluid, and victims are added and removed frequently.
- Extortion success is another key factor, if the amount of paying victims greatly increases, 'total' ransomware numbers may also appear to decrease.

With this said, we can draw some insight from this data with sensible assumptions – and recognizing the data isn't perfect, it does provide a decent gauge on the ransomware landscape. The assumptions the industry typically abide by are:

- There is a roughly relatively consistent month-on-month victim payment rate,
- Actor posts do contain an element of truth.

January's numbers fell very slightly from December's count of 516, to 508. However, while this is a drop from December, it is a year-on-year growth of 139 compared to January 2024 and is the 5th highest month when compared to a rolling 12-month period. Considering that January is typically one of the least active months this could bode badly for the year ahead.

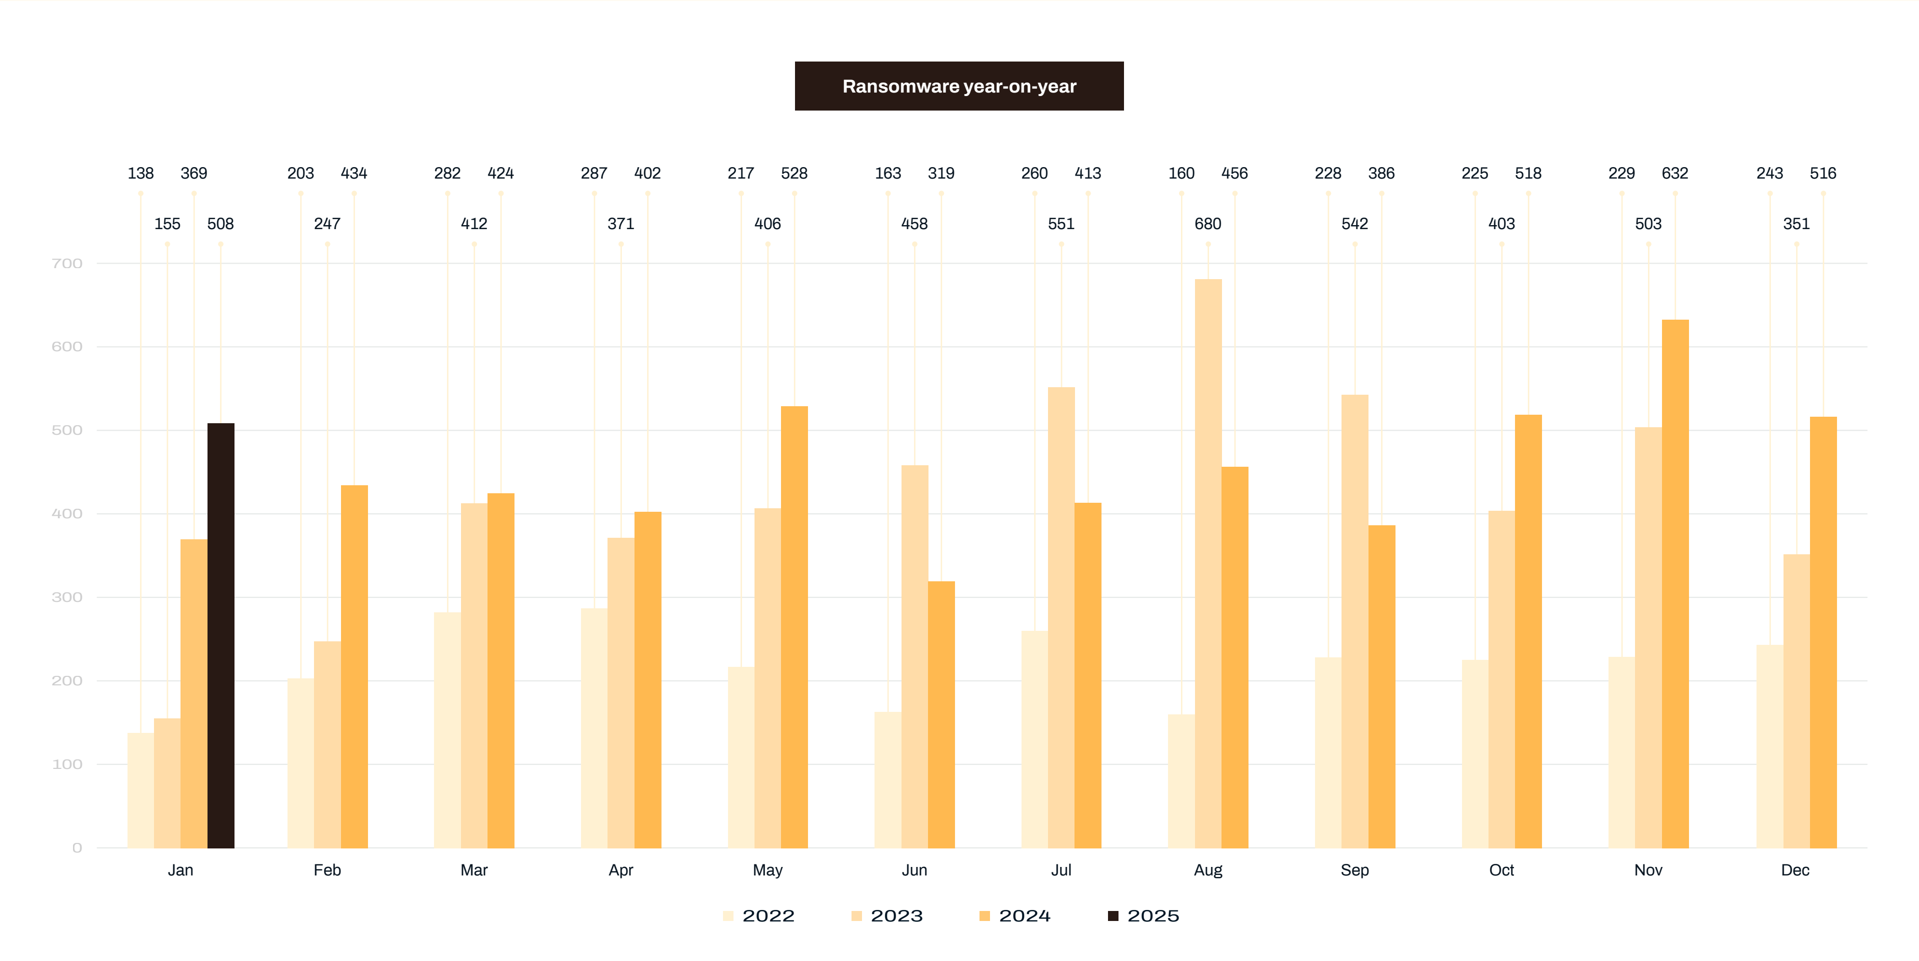
Ransomware year-on-year

The most active groups in January were Cl0p, who posted 63 victims, up 60 from their very quiet December, RansomHub, who posted 56 (down 13), Lynx who posted 42 (up 27), and Akira who posted 39 (down from 50).

The three groups that saw the biggest month on month increase were Cl0p, Lynx and INC Ransom, who cumulatively posted an extra 107 victims this month compared to last.

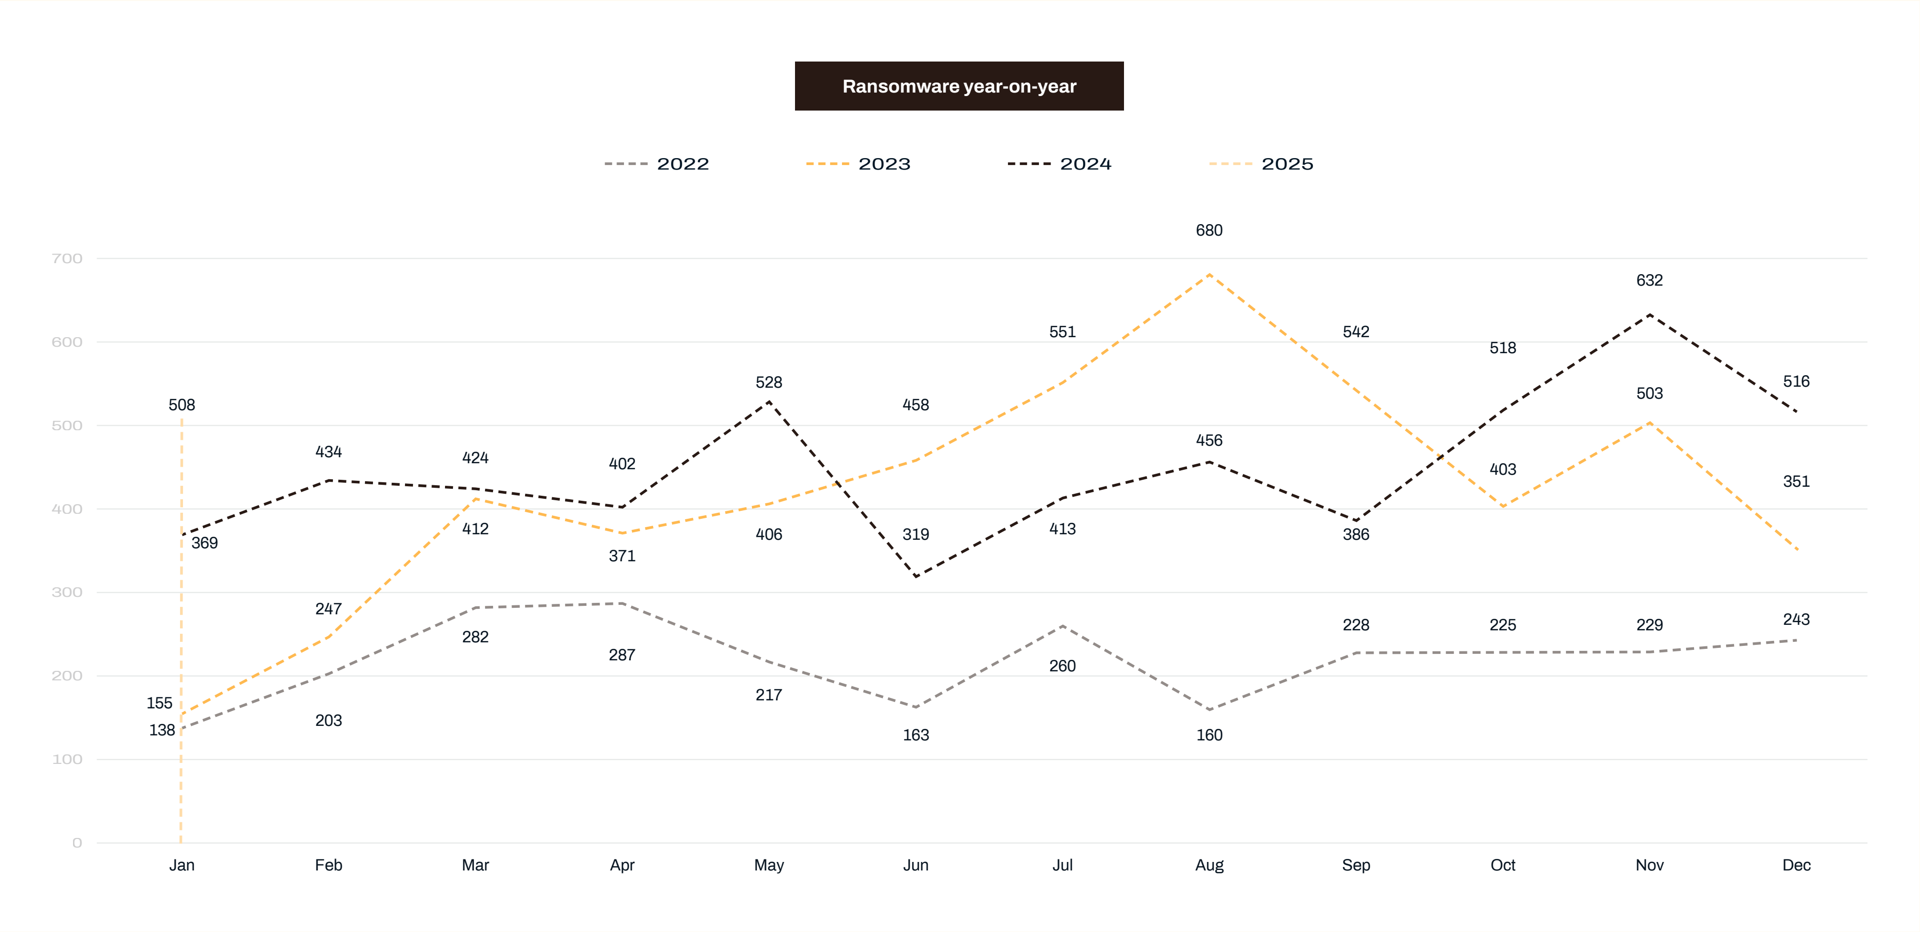

The most active groups in January were Cl0p, who posted 63 victims, up 60 from their very quiet December, RansomHub, who posted 56 (down 13), Lynx who posted 42 (up 27), and Akira who posted 39 (down from 50).

The three groups that saw the biggest month on month increase were Cl0p, Lynx and INC Ransom, who cumulatively posted an extra 107 victims this month compared to last.

Group	Dec-24	Jan-25	Change
CL0P	**3**	**63**	**60**
Lynx	**15**	**42**	**27**
INC Ransom	8	28	20
8BASE	11	25	14
Medusa	13	22	9
Space Bears	3	11	8
Monti	1	8	7
Safepay	14	21	7
Cicada3301	0	5	5
LockBit	5	10	5
Cactus	5	9	4
DragonForce	9	13	4
FSOCIETY	1	5	4
Qilin	20	24	4
Rhysida	5	9	4
3AM	2	3	1
Apos Security	0	1	1
Embargo	0	1	1
Everest	6	7	1
MedusaLocker	0	1	1

New groups

There were no new data leak sites in January2024.

FunkSec

CheckPoint have reported on a ransomware group named FunkSec who appear to have used LLM generated code, or at least LLM assisted code generation, to create their malware. FunkSec appear to be based in Algeria based on data from VirusTotal and references in their Ransom notes, and while their English language communication is quite basic in most situations, their malware, and even their encryptor, have verbose, clear, and well written comments which describe what each segment of code is intended to do, exactly as is seen in LLM generated/assisted code.

AI

Meta Llama framework vulnerable to data deserialization RCE

A large part of creating and using generative AI is data serialization and deserialization, something which software has been doing for a long time now. Both secure and insecure methods of doing this have been identified, and researchers have discovered that Llama uses an insecure method known as pickle. However, because LLMs are complex, high-level pieces of software with complex webs of dependencies, what is actually happening is that a software package that Llama relies upon uses the pickle format, something which the developers of Llama were unaware of. Unfortunately, it is quite reasonable that they would not be aware of this, as in modern software development it almost certainly isn't possible for a developer to perform a security review of every single line of code which they import or rely upon for their own code to function.

WithSecure Insight

This specific vulnerability was addresses in October, and is only being reported now, however this highlights that not only do LLMs face their own specific security challenges such as jailbreaking and prompt engineering, they also face all the standard software, cloud, and hardware security issues that non-LLM technologies face.

Yet another jailbreak method

Researchers at Unit42 have identified a new way to jailbreak LLMs by abusing their ability to evaluate text. The Likert scale is a rating scale used to measure the respondent's agreement or disagreement with a statement. It is often used in surveys for example, to ask "how strongly do you agree or disagree with the following statement…"

The researchers found that they could ask an LLM to use a custom Likert scale to score a statement, and then ask the LLM to generate example responses that align with the various points on the scale. The example the researchers give was a scale where a score 1 statement does not contain instructions on how to construct a bomb, while a score 2 statement contains detailed step by step information on constructing a bomb. They then simply asked the LLM for an example that would score a 1, and an example that would score a 2. It was found that this increased the success rate of jail break attacks by 60%, although due to the number of different LLM models currently available, they were not able to exhaustively test each one.

WithSecure Insight

It does not seem that there will ever come a time that all possible jailbreak methods are identified and defended against. Much like with standard IT security, it is likely that LLM security will be a constantly moving target.

Yet another jailbreak method

Researchers at Unit42 have identified a new way to jailbreak LLMs by abusing their ability to evaluate text. The Likert scale is a rating scale used to measure the respondent's agreement or disagreement with a statement. It is often used in surveys for example, to ask "how strongly do you agree or disagree with the following statement…"

The researchers found that they could ask an LLM to use a custom Likert scale to score a statement, and then ask the LLM to generate example responses that align with the various points on the scale. The example the researchers give was a scale where a score 1 statement does not contain instructions on how to construct a bomb, while a score 2 statement contains detailed step by step information on constructing a bomb. They then simply asked the LLM for an example that would score a 1, and an example that would score a 2. It was found that this increased the success rate of jail break attacks by 60%, although due to the number of different LLM models currently available, they were not able to exhaustively test each one.

WithSecure Insight

It does not seem that there will ever come a time that all possible jailbreak methods are identified and defended against. Much like with standard IT security, it is likely that LLM security will be a constantly moving target.

In Brief

Software supply chain

Fake NPM packages are being created and added to NPM which impersonate the Ethereum blockchain development environment, HardHat, with over a thousand known downloads of the malicious packages so far.

Multiple typo squatting NPM software packages targeting Solana blockchain users have been discovered which perform wallet draining attacks, sending stolen private keys to the attacker via Gmail SMTP servers.

The Windows version of a memecoin pump (as in pump-and-dump) tool, DogWifTools, was compromised to perform a wallet draining attack after the developers left a GitHub token for the repository exposed and publicly accessible in the source code. Losses are believed to be in the millions of dollars, and multiple releases of the tool were compromised before the attack was detected.

Identity

The Go language FastHTTP library is being used to perform high speed brute force attacks against M365 accounts with a reported 10% success rate. The user agent hasn't been changed, so it still shows as FastHTTP in logs. 65% of the traffic in these attempts originate from Brazil.

Opportunistic cloud and on-prem actor, TripleStrength monetize compromised clouds through cryptojacking and selling access/credentials, while monetizing on-premise networks through ransomware. Access seems to be through credentials compromised by Raccoon Stealer.

Sneaky, a Phishing-aaS kit is a full spectrum service for credential and 2fa prompt/theft in one. It uses compromised WordPress sites for hosting and creates fake authentication pages to harvest credentials and 2fa codes.

DDoS

A Mirai botnet, first discovered in February 2024 has begun exploiting vulnerable Four-Faith industrial routers, with 15,000 active bot nodes per day, more than Cloudflare reported took part in the 5.6Tbps DDoS attack below.

Cloudflare observed a 5.6 Tbps DDoS attack lasting 80 seconds, launched from a botnet of 13,000 IoT devices (not a particularly large number of devices considering modern smart device security and exploitation trends!)

Researchers discover that it is possible to cause OpenAI's internet crawler to DDoS a destination site, as URLs for it to scan be submitted in bulk lists, with no upper limit per post request, and no deduplication is applied to the submitted URLs. OpenAI and Microsoft have declined to acknowledge or address the report.

Other news in brief

SonicWall Appliance Management Console CVE-2025-23006, an unauth-RCE is being exploited by attackers. These are management devices, much like FortiManagers, which also recently had a high severity vulnerability exploited.

FortiGuard CVE-2025-55591 is believed to be the unknown zero-day that ArcticWolf reported was being exploited in November 2024. It is a 9.6 Authentication bypass vulnerability affecting FortiOS and FortiProxy management interfaces, allowing a remote attacker to gain super-admin privileges.

Over $85 million dollars was stolen from the Phemex crypto exchange. Initially only $29 million was reported stolen, but over several days it was gradually raised to $85 million.

Backdoor discovered in Chinese made patient health monitors which would send all patient data to a hardcoded IP address belonging to a Chinese university. CISA contacted the manufacturer asking for remediated firmware, and multiple times they sent back lightly modified firmware that still contained the backdoor.

International law enforcement operations seized multiple cybercrime forums which are involved in the trade of stolen credentials, and credential stealing tools and techniques.

The full impact of Salt Typhoon's compromise of the Western Telecommunications industry is unknown, but it is believed to be substantial and international. The full impact may not be known however, as incoming US President Trump has fired the entire CISA Cyber Safety Review Board team who were investigating. It is presently unclear how this will improve cyber security outcomes.

Threat data highlights

Exploits

WithSecure vulnerability exploit detections rose hugely this month. While some of that change is probably attributable to how quiet December tends to be, there does seem to have been something else going on. The largest increase detection was for a 2021 Linux Kernel RCE, which increase from 7 detections to 591 detections, a roughly 8,000% increase. The rest of the top 5 increases were all Java exploits, with the youngest of those being 11 years old.

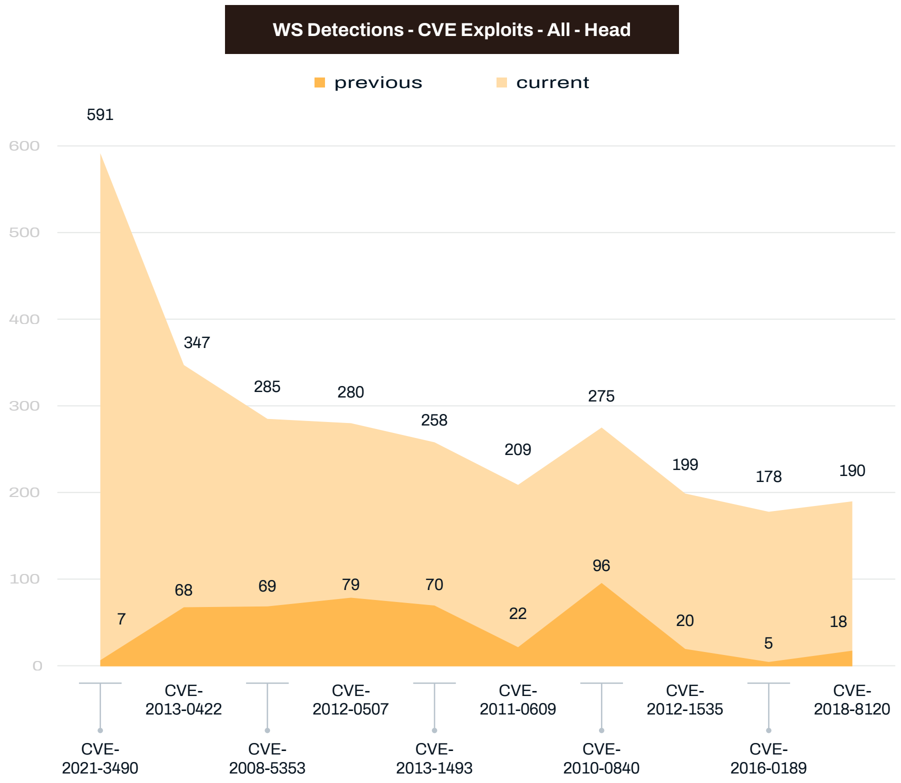

There are no decreases of interest in W/S detections this month.

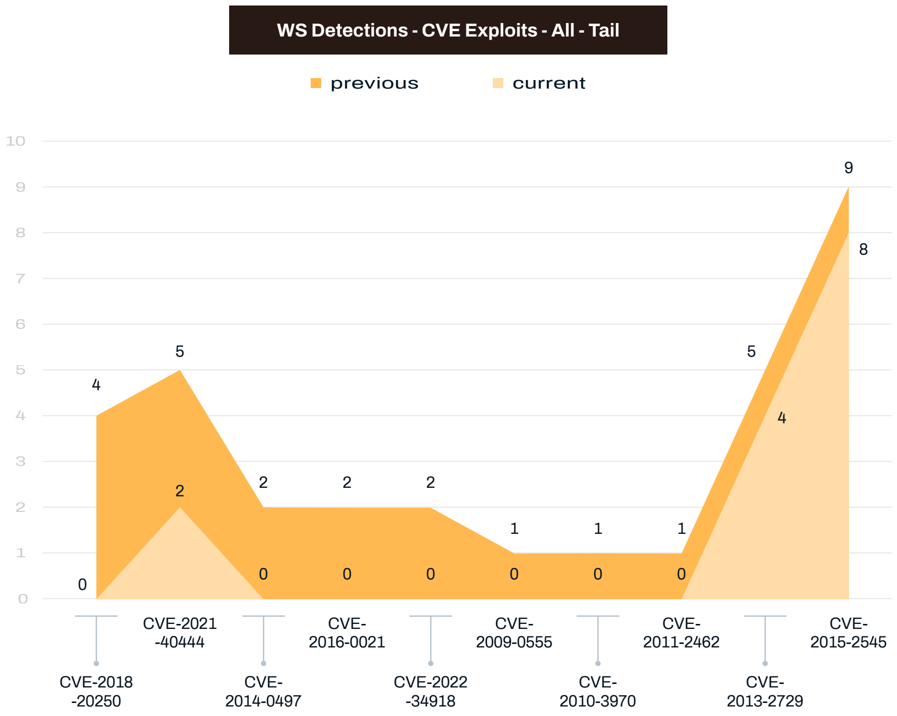

In detections which are new this month compared to last month there are two which have shown quite large increases, they are a 2018 Foxit PDR reader RCE, and a Windows <8.1 privesc to kernel vulnerability:

WS Detections - New - Monthly

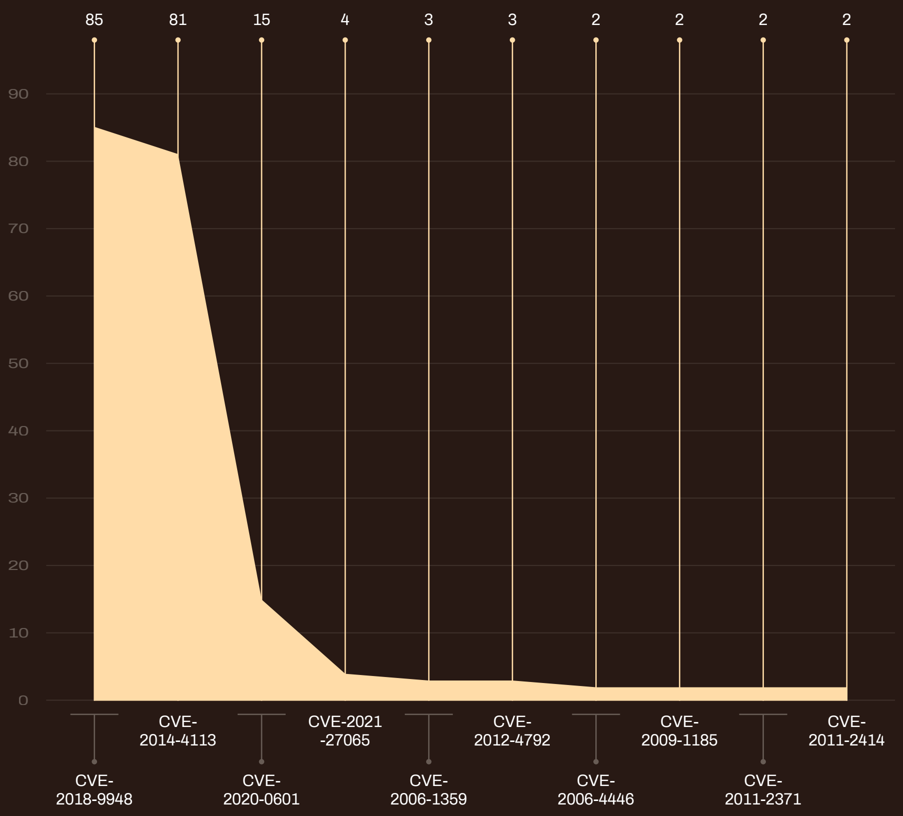

There are no particular changes of interest in VirusTotal 2023/24 CVE detection increases.

Virus Total 2023/24 CVE Exploit Detections - First submitted 2023/24 - Head

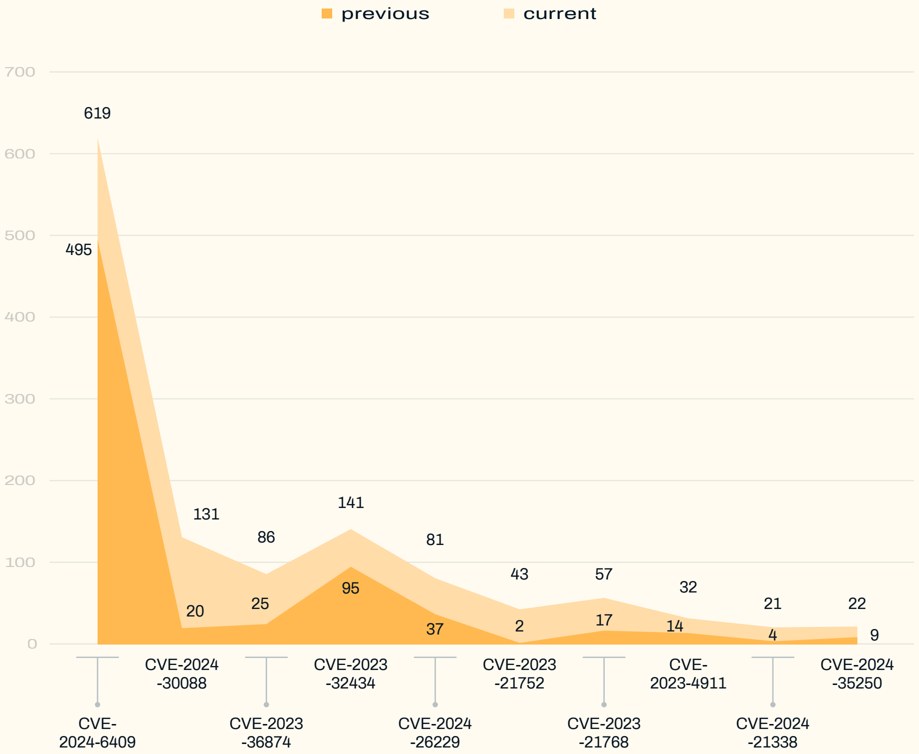

Looking at the VirusTotal detection decreases there was a relatively large drop in detections of the Copy2Pwn Windows Mark of the web bypass this month. What makes this particularly interesting is that a new Windows Mark of the web bypass in 7-zip was disclosed this month. While a patch has been available since the end of November, this vulnerability is known to have been exploited by Russian actors targeting Ukraine.

Virus Total 2023/24 CVE Exploit Detections - First submitted 2023/24 - Tail

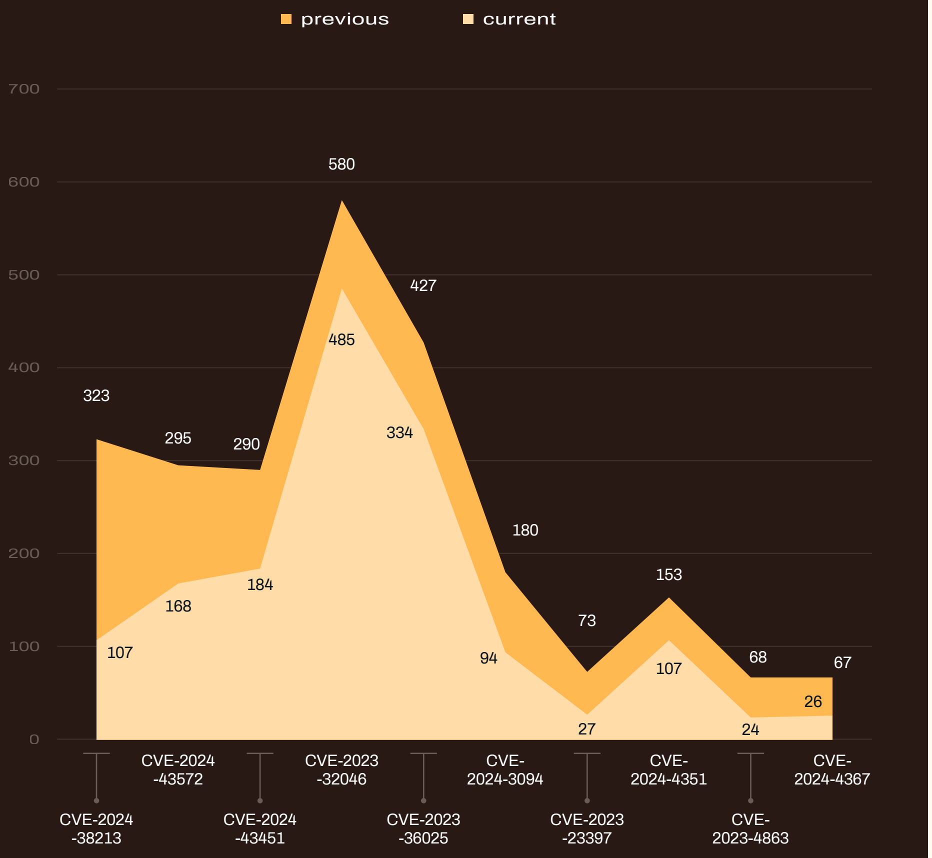

In VirusTotal detections of all time CVE exploits, there are significant increases in a number of vulnerabilities. At 1st and 2nd, the duplicate TCP/IP RC4 and Checkpoint Firewall IKE exploit detections have doubled to over 7,000, while at 3rd place detections of an IE 9-11 cross domain/zone information disclosure vulnerability have increased 10-fold to over 3,600. Further down the graph, at 7th and 8th place there appear to be 2 more duplicate vulnerabilities which have increased from 926 to 2814, both of which are DNS poisoning attacks. At 10th place a 2016 DES/TripleDES birthday attack has also increased drastically, from 693 to 2140.

Virus Total CVE Exploit Detections - First submitted 2023/24 - Head

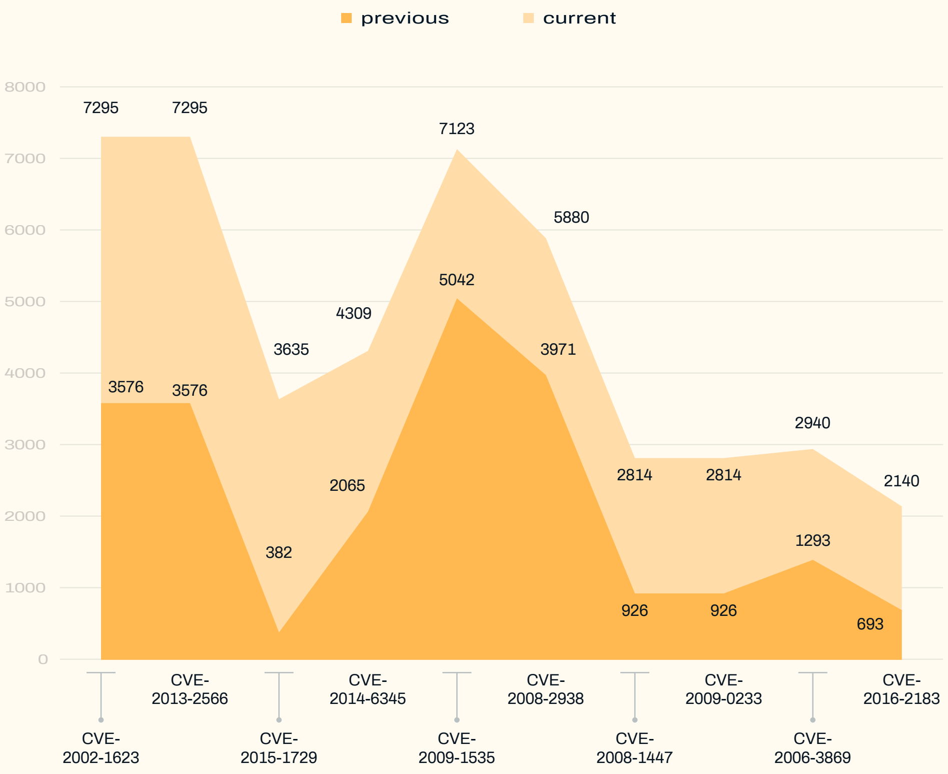

Finally, in VirusTotal all time CVE exploit detection decreases, we can see that the same 4 DoS vulnerabilities are still taking up the top 4 slots, while the 2018 Dasan GPON router vulnerability which we discussed before has also dropped significantly at 6th place, and the 2017 Huawei HG532 router vulnerability is also still dropping at 9th place. Both the GPON and Huawei vulnerabilities would be very useful to anybody building a SOHO router botnet.

Virus Total CVE Exploit Detections - First submitted 2023/24 - Tail

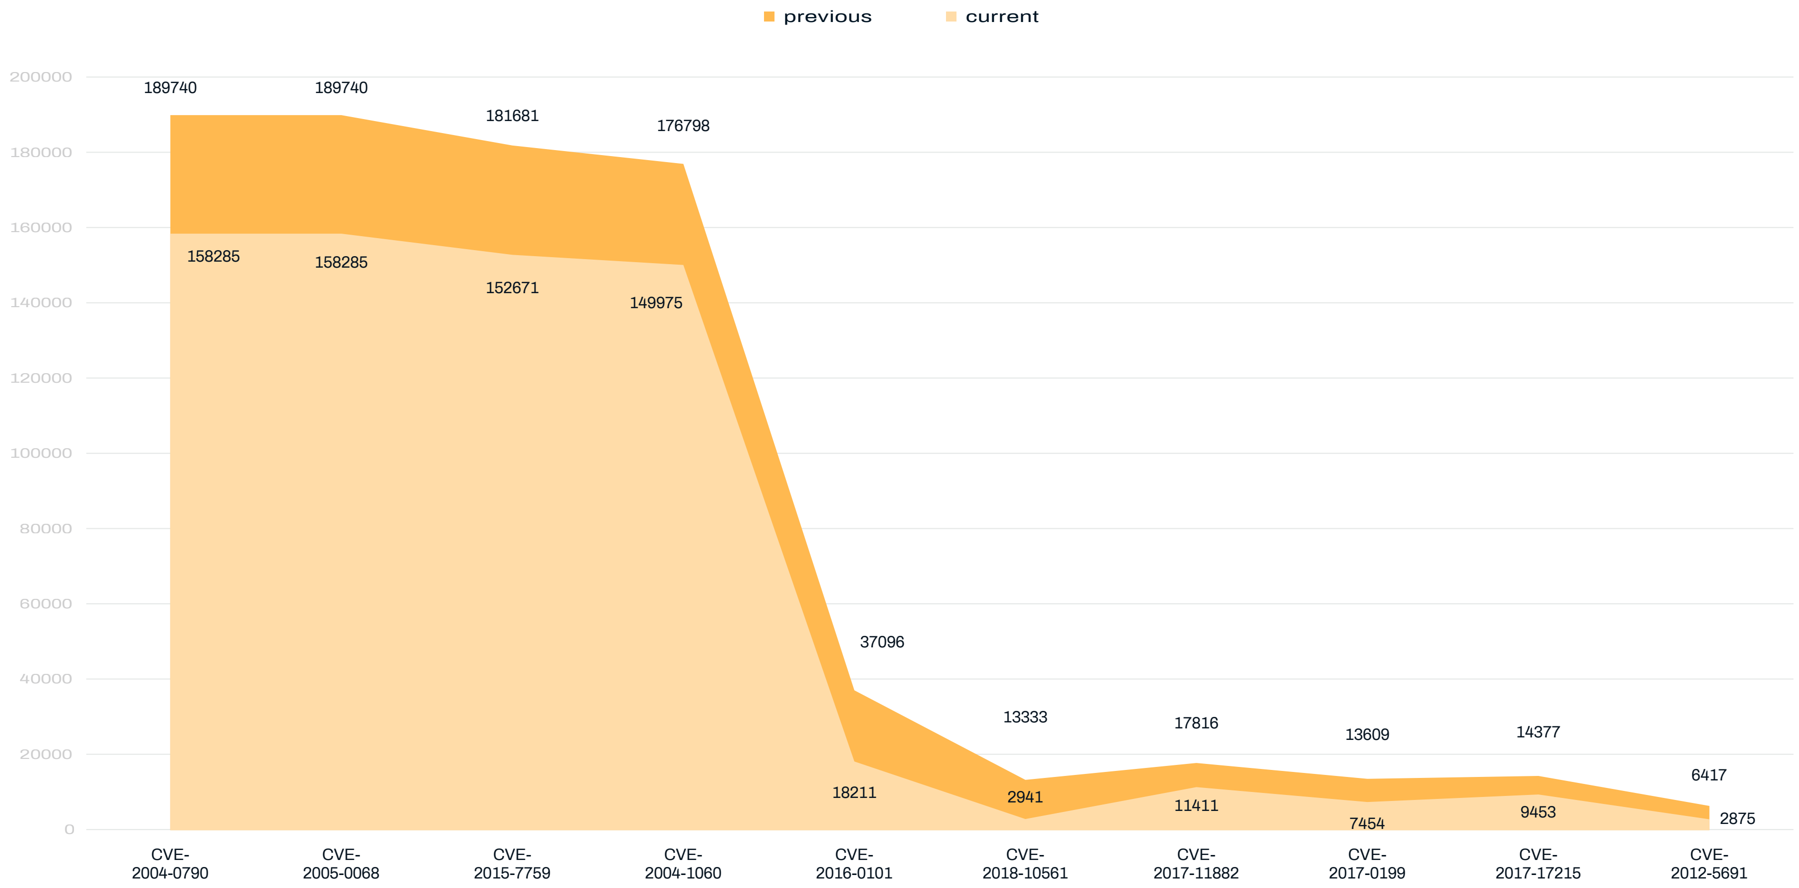

Phishing malware delivery

In phishing malware delivery stats this month, Loki malware is by far the highest volume of malware observed delivered via phishing, with 20 times the number of detections of any other malware.

Malware	Total count
loki	13956
formbook	703
agenttesla	648
strelastealer	599
snakekeylogger	455
dbatloader	308
asyncrat	156
guloader	113
masslogger	110
strrat	88
remcosrat	56
quasarrat	50
redlinestealer	33
rustystealer	1

Loki showed large spikes at the end of the month in North and South America, Africa, Oceania, and Europe.

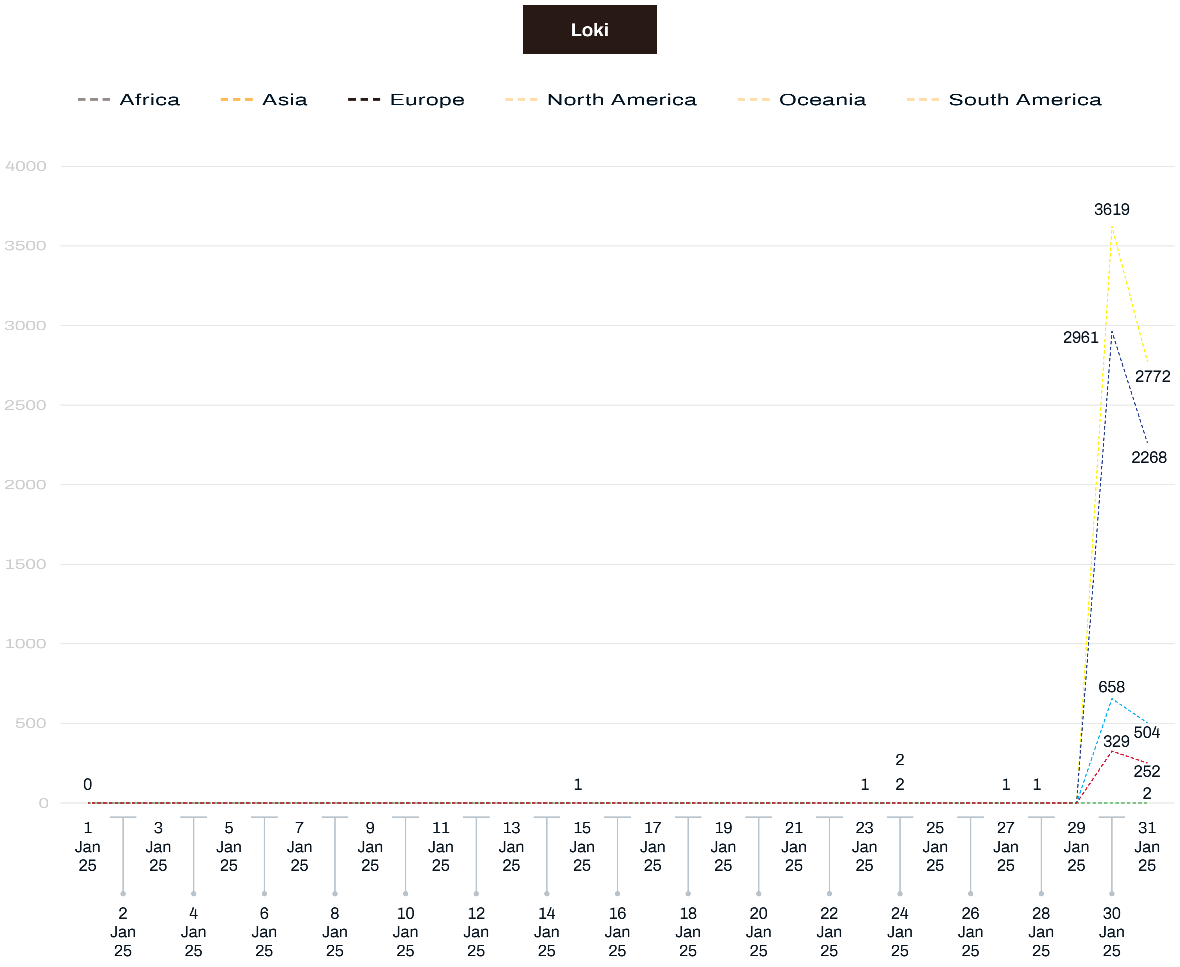

Formbook was the next highest malware internationally, remaining active throughout the month. Although it was most seen in Europe, this may of course be an artifact of our heavier presence in Europe.

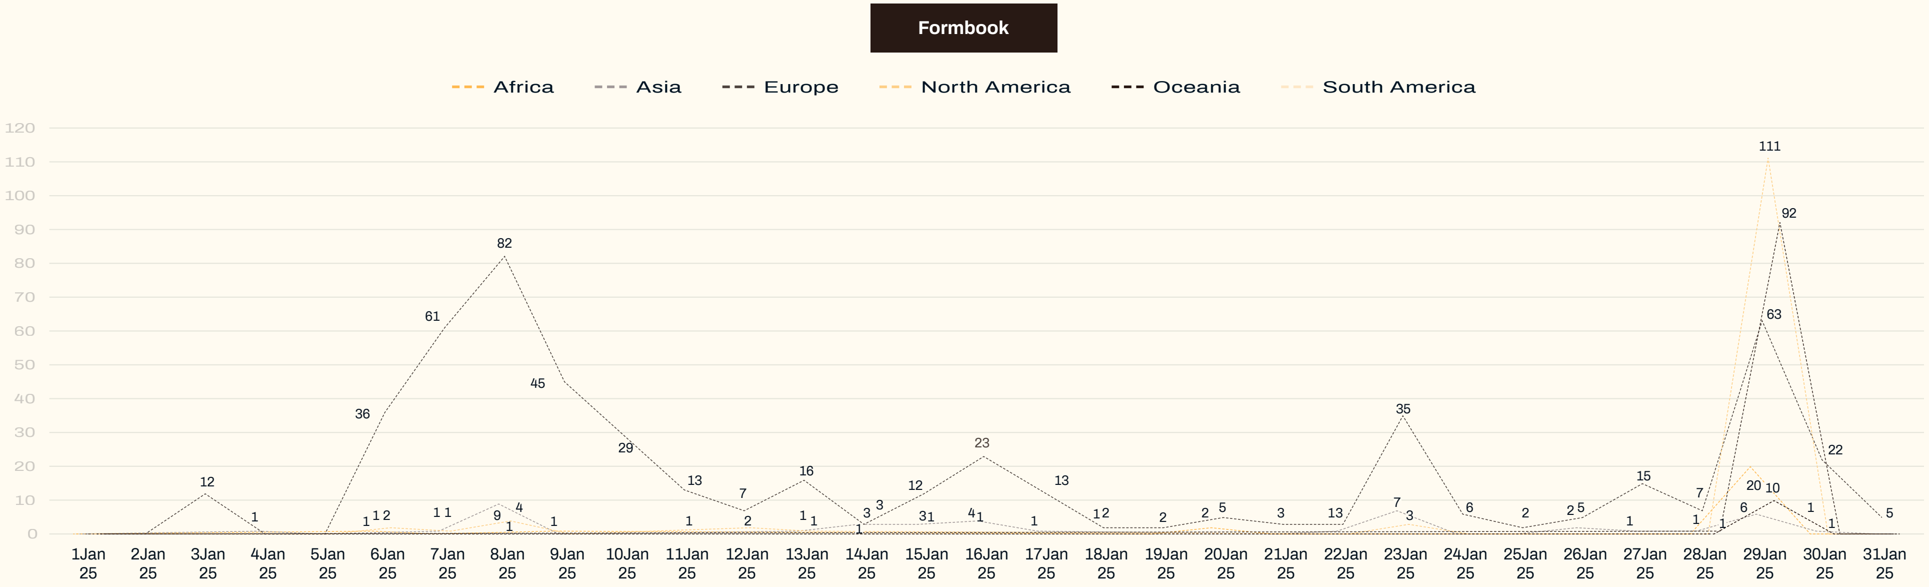

Significant spikes were seen in Europe in the week beginning 20th of January for StrelaStealer, AgentTesla, DBatLoader, and ASyncRat, which were not seen elsewhere.

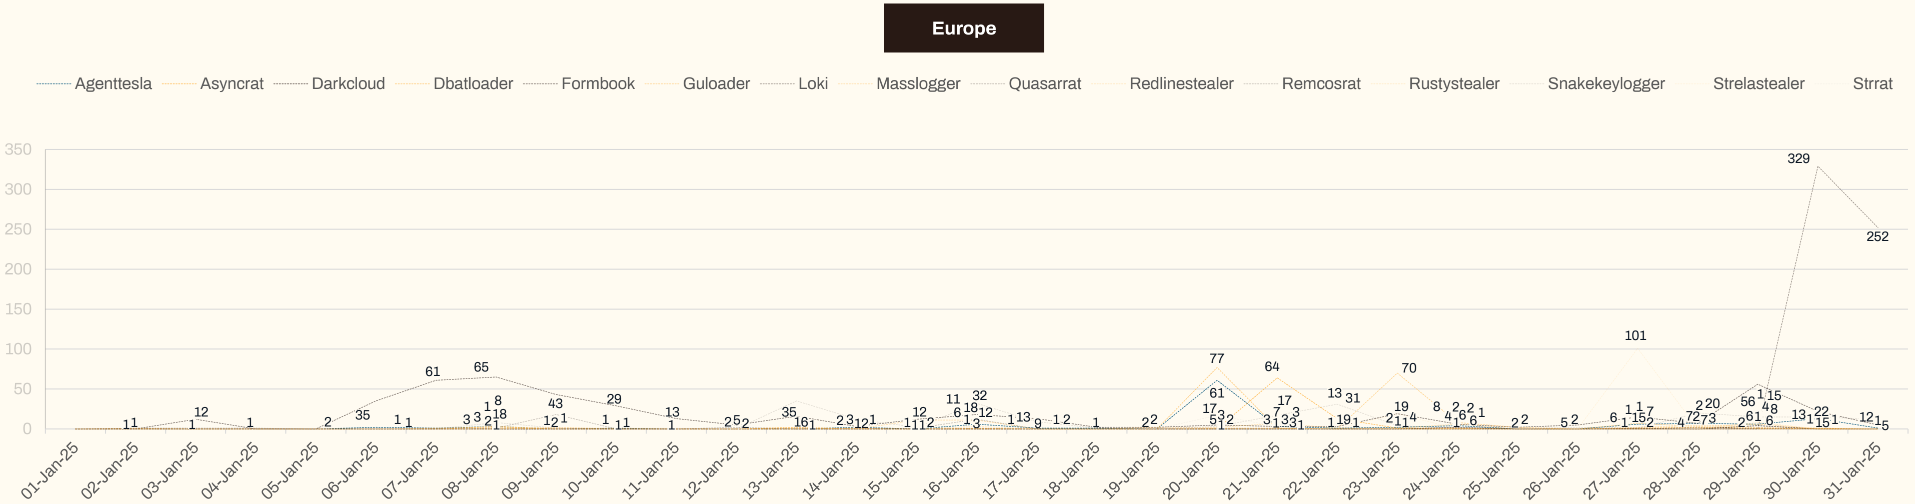

About WithSecure™

WithSecure™, formerly F-Secure Business, is Europe's cyber security partner of choice. Trusted by IT service providers, MSSPs, and businesses worldwide, we deliver outcome-based cyber security solutions that protect mid-market companies. Committed to the European Way of data protection, WithSecure™ prioritizes privacy, data sovereignty, and regulatory compliance.

Boasting more than 35 years of industry experience, WithSecure™ has designed its portfolio to navigate the paradigm shift from reactive to proactive cyber security. In alignment with its commitment to collaborative growth, WithSecure™ offers partners flexible commercial models, ensuring mutual success across the dynamic cyber security landscape.

Central to WithSecure's™ cutting-edge offerings is Elements Cloud, which seamlessly integrates AI-powered technologies, human expertise, and co-security services. Further, it empowers mid-market customers with modular capabilities spanning endpoint and cloud protection, threat detection and response, and exposure management.

WithSecure™ Corporation was founded in 1988, and is listed on the NASDAQ OMX Helsinki Ltd.