Threat Highlight Report

April 2024

WITH secure

Contents

Foreword

This month we have been given a very clear statement on the true cost of a ransomware attack compared to the cost of the ransom. Change Healthcare have published their Q1 financials which estimate a cost of over $1.6 Billion from an attack that resulted in a ransom of $22 million being paid. We also had a clear illustration of just how difficult security can be for even the most knowledgeable of organizations, as both CISA and MITRE disclosed breaches due to Ivanti ConnectSecure. In the case of MITRE this was done with an admirable level of transparency from which others can benefit. We also report on the results of an investigation by CISA into the Microsoft Exchange Online compromise of late 2023, which comes out with some rather scathing observations and recommendations regarding the Microsoft corporation's culture, risk management, and communications.

We have a quick run through a number of news items in brief for the month, and an analysis of exploit detection data, which includes an interesting spike in detections of files targeting an old but unpatchable Huawei router vulnerability. We also have newly published WithSecure research which includes the excellent new Kapeka malware report, which has formed the leading edge of a wave of analysis of the Russian Sandworm APT group (designated APT44 this month) from across the cybersecurity industry.

- Stephen Robinson, Senior Threat Intelligence Analyst, WithSecure

1 Monthly highlights

1.1 Change Healthcare receive second ransom demand

As happens at this time of year, United Healthcare, owners of Change Healthcare, have published their Q1 results, which includes the costs related to the ransomware incident that completely disrupted their operations. They state that in Q1 it costs them $872 million, with a predicted total cost by the time it is fully resolved of $1.6 billion. That is almost 1% of the entire UK National Health Service budget for one year, or roughly 25 times the losses due to ransomware reported to the FBI in 2023. As well as these direct costs, Change Healthcare have also made interest free loans and advances totaling $6billion to customers. In a survey by the American Medical Association, four out of five respondents said that they had experienced direct financial impact due to the incident, with 48% of respondents entering into new agreements with service providers other than Change Healthcare.

Meanwhile on the other side of the fence, Change Healthcare were listed on the RansomHub leak site with the exact same volume of data as was stolen in the ALPHV attack. RansomHub are a new group we discussed in the March THR who have upended the usual payment model for RaaS brands and are allowing affiliates to directly receive payments from victims. Ransomhub then relies on the affiliates to forward on a share of the payment to the central brand. We theorized that this could have been a direct response to the ALPHV exit scam, which saw ALPHV simply take the entirety of the ransom for themselves. Now, while it cannot be proven at present, it does appear that the affiliate who hacked Change Healthcare under the ALPHV brand and who stated that they still had the stolen data, has now signed up with RansomHub to re-ransom Change Healthcare. A few days before the Ransom deadline for leaking the data, Change Healthcare were removed from the leak site. This suggests that they may have paid a second ransom to prevent the data being leaked. It may seem unlikely that an organization would pay multiple ransom demands like this, but the information available suggests that a second $22 million dollar ransom payment would hardly register, being only 0.014% of their estimated total costs from the incident. As an additional indicator of just how well motivated Change Healthcare would be to avoid a data leak, they have stated that the stolen data likely contained the PII of a substantial proportion of all US citizens.

Finally, blockchain researchers have been watching the ALPHV cryptocurrency wallet which received the original $22million ransom payment closely, and have finally started to observe the money being moved out of the wallet and laundered via cryptocurrency mixers and exchanges.

WithSecure Insight

While the initial ransom demand, and probably also the second demand were for tens of millions of dollars, we now know that this is just a drop in the ocean compared to the total cost to Change Healthcare. In fact, as if we needed any more sobering numbers to take from this incident, we can see that the cost of the ransomware incident was multiple orders of magnitude higher than the cost of the ransom itself. As ever, our advice is that you cannot trust cyber criminals, and that it is far cheaper to prevent security incidents than it is to recover from them.

1.2 MITRE and CISA disclose breaches due to Ivanti ConnectSecure

The impact and depth of the ongoing security issues surrounding Ivanti ConnectSecure (ICS) have been well discussed, and just this month the CEO of Ivanti released a 6 minute video stating that in response to this security incident the company would begin implementing a 'Secure By Design' ethos for their security products. This is a very positive move and shows real bravery by risking criticism from those who might raise concerns as to what Ivanti's design ethos was before this incident. We are optimistic that this will raise the bar for threat actors to utilize security products as infection vectors to gain access to victims. Just this month two notable victims have come forward which give a perspective into just how difficult this incident was for Ivanti's customers to defend against, even for those staffed with cybersecurity experts.

CISA, the US Federal Cybersecurity and Infrastructure Security Agency disclosed that the Chemical Security Assessment Tool (CSAT), a security tool/service which they provide to help secure chemical plants, was compromised through a vulnerable ICS instance. This led to a data theft affecting more than 100,000 individuals, all of whom are very likely to be closely involved in the security of critical national infrastructure. CISA stated that they implemented vendor-recommended fixes for the known ICS vulnerabilities on January 11th and ran daily checks with the Ivanti Integrity Checking Tool (ICT) after that. On January 26th they discovered that CSAT had been compromised, and that attackers had had access to the device for two days. The attackers were able to bypass both the Ivanti specified mitigations, and the Ivanti ICT checks. This may explain why CISA made a public statement (covered in last month's THR) that the Ivanti ICT was not sufficient to check for compromises, though Ivanti have released multiple versions of the ICT as the incident has progressed, as well as changing their mitigation advice.

MITRE, the research organization who maintain the MITRE ATT&CK cybersecurity tactic and technique matrix also disclosed that they were compromised via an ICS appliance. The attackers compromised an unclassified research and development network in early January, prior to the initial disclosure of ICS zero-days. While MITRE patched the ICS appliance, the attacker had already moved laterally into their VMWare infrastructure, compromising an administrator account then deploying backdoors and webshells to maintain persistence and harvest further credentials. MITRE have published details of the compromise and their response, and have stated that they will share as much information about the incident as possible to allow others to learn from it.

WithSecure Insight

It is an overused term, but there truly was a perfect storm around Ivanti Connect Secure this year. These zero-day vulnerabilities were hugely impactful and easily exploited en masse. Organizations as security aware as MITRE and CISA suffered compromises, even though they did all the right things and implemented the recommended responses as soon as they could. CISA were actively watching for compromise, and it still took two days to identify. MITRE patched and mitigated the vulnerable device, but by that point the attacker had already moved laterally from there and deployed multiple persistence methods. It is really quite chilling to consider just how many organizations were vulnerable to attacks that were able to compromise even entities such as these. We can only hope that every organization whose purpose is to reduce threat surface, and not become part of it (and we certainly include ourselves in that list) take some hard earnt and valuable lessons to heart from this incident.

1.3 Review of 2023 Exchange Online compromise heavily criticizes Microsoft security

The compromise of Microsoft's Exchange Online email server in mid-2023 had a significant impact on the US government. As such, CISA's Cyber Safety Review Board (CSRB) have conducted a review into the incident and concluded that the incident was preventable and should not have occurred. The executive summary of the review states that:

- "Storm-0558 was able to succeed because of a cascade of security failures at Microsoft"
- "Microsoft's security culture was inadequate and requires an overhaul"
- "The Board identified a series of Microsoft operational and strategic decisions that collectively point to a corporate culture that deprioritized both enterprise security investments and rigorous risk management"

The review also criticized Microsoft's communication, as they inaccurately stated that they had determined the likely root cause of the incident, when in fact they did not and still do not know the root cause. The review even goes so far as to suggest that Microsoft should de-prioritize feature development across their cloud infrastructure and product suite until such a time as substantial security improvements have been made, which gives us a real impression of the seriousness and severity of the findings of the report.

It is scathing, but it is also an interesting read, and makes a good argument that having almost unlimited resources to throw at a problem does not guarantee success. To illustrate this, the report states that one week after they were notified by a US Government customer of the intrusion, Microsoft realized that the attacker must have acquired a specific type of valid Microsoft cryptographic certificate. Microsoft assigned this their highest urgency level and engaged multiple teams to investigate, developing 46 separate hypotheses as to how the key may have been acquired. These hypothetical scenarios included that the adversary possessed a theoretical quantum computer able to break public-key cryptography, or that an insider threat had stolen the key during its creation 7 years earlier. Microsoft then assigned teams to investigate each of these hypotheses, and after 9 months has stated that these investigations are still ongoing.

WithSecure Insight

It would be very difficult to criticize MITRE or CISA for the compromises they experienced. While it might be easier to criticize Microsoft for this and other similar compromises during 2023, we do not believe such harsh judgement may necessarily be warranted. Microsoft is a multinational company with thirty-fivefifty times as many employees as CISA. We often say that security needs to be considered and implemented from the beginning, as it is far more difficult to implement afterwards. This is especially true in complex environments. Microsoft is a huge and complex organization which could be compared to an oil tanker, having to make decisions and implement course changes miles in advance. As such, implementing the organizational security improvements and prioritizations recommended by CISA will likely take significant effort over some time.

It remains to be seen the extent to which Microsoft can successfully make these security course corrections, but considering just how heavily the information technology industry relies upon them, it will have (hopefully positive) implications for organizations of all sizes.

2 Ransomware: Trends and notable reports

2.1 The numbers

Numbers slightly dropped from last month's total (424) to 402. This represents an 8.36% increase from the same month in 2023, and a 40.03% increase from the same month in 2022. While numbers decreasing for the second month in a row is good news, over the last three years the second quarter of the year has presented a decrease in total victims.

Ransomware	Count	Change
3AM	1	=
8BASE	25	+8
Abyss	1	-4
Akira	18	-3
Alphv (BlackCat)	0	-8
Apos Security	4	+4
BianLian	11	-6
BiteMe	0	-1
BlackBasta	24	-11
Blackout	1	+1
Blackbyte	0	-1
Blacksuit	21	+13
Cactus	13	+3
CHCC Leak	0	-1
CiphBit	4	+4
Cloak	8	=
CL0P	3	-1
Daixin	1	+1
dAnon	8	+8
DarkVault	17	+17
Data Leak	9	+8
Defray777	1	-4
Dispossessor	1	-4
Donex	0	-5
Donut Leaks	0	-2
DragonForce	12	+6
Dunghill Leak (News)	1	+1
Embargo	2	+2
Eraleignews	2	+2
Everest	2	-1
HelloGookie	0	-4
Hunters International	29	+11
INC Ransom	15	+2
Kill Security	1	-4
LockBit	25	-24
MalekTeam	2	+2
Mallox	1	-2
Medusa	28	+1
Meow	0	-2
Mogilevich	0	-4
MyData	2	+1
Play	30	-16
Qilin	12	1
Qiulong	6	+6
RA Group	14	+3
Ransomhouse	6	+6
RansomHub	23	+1
Red Ransomware	1	-11
Rhysida	6	+2
Snatch	0	-9
Space Bears	8	+8
Stormous	1	-10
Trigona	0	-6
Underground	2	=

Monitoring Lockbit

Lockbit leak site numbers have halved for the second month in a row, dropping from 107 (February) to 49 (March) to 25 in April. This is a positive sign that trust has significantly eroded in the brand, however the industry must still be open to the possibility that key Lockbit associates are working to reinvent the program.

For the first time, Play ransomware takes the top spot this month, despite a reduction of 16 victims. Hunters International, 8Base and BlackSuit also posted a notable increase. Medusa's increase in victims posted coincided with the Lockbit LEA action, and this has continued into this period.

New groups

There are seven newly observed leak sites that have begun posting victims this month:

- **Qiulong** posted six victims in April. All companies posted were headquartered in Brazil, indicating a specific targeting profile. Furthermore, all bar one victim were small medical practices.
- **dAnon:** dAnon posted eight victims this period. There are no discernible patterns in victimology.
- **DarkVault:** DarkVault posted the most victims of all the newcomers this period with 17. There was a large geographical distribution of victims, with organizations in the US, UK, Belarus, India and Saudi Arabia.
- **Embargo:** A new leak site, posting only two victims in April.
- **Eralignews:** Details of Eralignews emerged when independent researcher Rakesh Krishnan documented the ransomware in a blog post. It appears that the actor has attempted to assign itself the designation APT73. Numerical APT (Advanced Persistent Threat) references are a schema first used by Mandiant, reserved for advanced and persistent threats. Often to define capable nation-state intrusion sets, this ransomware intrusion set is a long way away from being considered an APT. Krishnan notes the actor's Twitter/X account follows a number of registered Finnish accounts, suggesting there is a connection to Finland with the actor.
- **Space Bears:** Space Bears posted 8 victims this period and is believed to be associated with the leak site of Phobos ransomware, a prominent ransomware family related to the 8base leak site. While there are different variants of Phobos ransomware, there is no ransomware leak site branded specifically to Phobos. A post this month from Microsoft's Threat Intelligence team stated Phobos was among the most prominent variants used in Q1 2024.
- **Apos Security:** Apos Security posted four victims in April, two of them in Brazil.

2.2 HelloKitty rebrands, say HelloGookie

Someone posting under the name Gookee/Kapuchino and claiming to be the original creator of the HelloKitty RaaS brand has announced a relaunch/rebrand of the RaaS group to HelloGookie. At the same time, they also released passwords for previously leaked data, including leaked videogame source code from CD Projekt Red and Cisco network information. The source code and locker builder for HelloKitty was leaked by a user going by the same Gookee/Kapuchino name in November 2023, and in leaked private conversations the HelloKitty developer used the name Guki, so there is indeed reason to believe these actions have all been carried out by the same person/group.

2.3 CISA say Akira has stolen over $40million from 250 victims

CISA published an advisory containing TTPs of the Akira ransomware brand, suggesting that their operations have reached a point that the US government feels they must begin to address the issue. The statistics published within the advisory show why this might be the case, stating that as of January 1st 2024 Akira had impacted over 250 organizations, claiming around $42 million in ransoms.

2.4 Large volume of data stolen from UN in ransomware attack

The United Nations Development Program issued a statement that they had been compromised by a ransomware actor who had stolen large amounts of information. This ties to a post on the 8Base leak site from late March which stated they had compromised the organization. There is no indication of how the compromise occurred, though in 2021 the UN acknowledged publicly that there had been multiple successful attacks against them as a result of stolen login credentials to internal UN projects being sold on the dark web.

2.5 1.3TB and control of local streetlights stolen from UK local council

The UK local government organization Leicester City Council was a victim of an INC Ransomware attack in March. As per UK government policy, they did not pay the ransom demand, and now 1.3TB of stolen data has been leaked by the attackers. In an interesting illustration of the breadth of the cyberattack it has been noticed that the streetlights in the city are now on permanently. A council response to a query about the streetlights states that due to shutting down systems after the ransomware attack the central management system for the streetlights is “misbehaving”, and some streetlights have gone into a fault mode which cannot yet be centrally resolved. In the event of a fault like this, the streetlights are configured to fail safely and stay permanently lit to ensure that no roads are instead left unlit. While this in particular is not a particularly severe ransomware symptom, it’s definitely unexpected, at a time when Leicester council will almost certainly be experiencing severe disruption to operations.

2.6 Peruvian hosting firm faces $140 million dollar ransom demand

In a compromise reminiscent of the recent Tieto Evry compromise by the Akira ransomware brand, the Chilean hosting provider IxMetro Powerhost has been compromised by the SEXi ransomware brand. IxMetro Powerhost are based in South America, but also active in North America and Europe. The SEXi brand, who are seemingly named after the VMWare ESXi hypervisor software, encrypted Powerhost’s ESXi servers and backups, causing extended outages for Powerhost and their customers. Powerhost’s CEO said in a statement that the attackers demanded 2 Bitcoins per customer for decryption, which comes to a total of $140million.

2.7 Statistics: Payments drop while mass exploitation surges

Researchers from Coveware have published an analysis of ransomware trends in 2024Q1, stating that they have observed a continuing decline in the number of ransoms being paid, with only 28% of ransom demands being met this quarter, compared to 29% in 2023Q4. They also observed a 32% drop in the average ransomware payment to $381,000 along with a 25% increase in the median payment to $250,000, which means that there are fewer very large payments, but that the typical payment amount has increased. Another interesting statistic they give is around initial vectors. The highest infection vector is once again “unknown”, however the second highest (and increasing) is remote access compromise, while phishing has seen a significant decline and now sits at the same level as Software Vulnerability.

Coveware state that the most common ransomware variants in 2024Q1 are Akira at 21% of attacks, followed by BlackBasta and Lockbit on 9% each. This means that Akira have been top of their rankings for 3 straight quarters. Interestingly, the “roll your own” non-RaaS Phobos ransomware takes joint third place along with Medusa, which may be a sign of the increasing levels of distrust among the RaaS industry after recent law enforcement takedowns and the ALPHV exit scam.

3 Other notable highlights in brief

3.1 Mandiant observe China and Russia increasingly targeting edge devices and infrastructure

Mandiant have published their annual report into IR cases that they were involved in in 2023, and they have observed a dramatic increase in Chinese and Russian state sponsored attackers targeting edge devices and infrastructure. They state that this activity was often enabled by exploiting zero-day vulnerabilities in those infrastructure devices, while avoiding deploying malware on Windows devices as much as possible. Mandiant put forwards the theory that this is an intentional choice to avoid EDR. This correlates with observations that we have made in previous THRs regarding infrastructure and edge service compromises. Echoing the recent statistics from Symantec (March THR) and Coveware, Mandiant saw compromise via exploitation increase 6% to 38%, while phishing dropped 22% to 17%. In contrast to Coveware, Mandiant only give statistics for compromises where the initial vector was identified. Do also note that Mandiant's statistics are for all intrusions investigated, while Symantec and Coveware only give statistics for ransomware.

WithSecure Insight

This is one of multiple reports recently that have observed an increase in Mass Exploitation of vulnerable edge services, a trend that WithSecure have also noted. With each additional report on the topic, each of which is using a different data set, the strength of this conclusion grows.

3.2 CrushFTP vulnerability under active exploitation

A zero-day vulnerability in CrushFTP was detected by Airbus's CERT which enables an attacker to escape their VFS (Virtual File System) and download system files. A CVE reference was not made available for this vulnerability for some days after the vulnerability was made public on the 19th April. Exploitation of this zero-day has been described as "possibly politically motivated", with actors undertaking targeting intrusion and intelligence gathering at multiple US entities.

WithSecure Insight

While we don't know who these entities are at present, we know that file transfer services are not just of interest to espionage actors. Targeting vulnerabilities in popular file transfer services is also popular for financially motivated threat actors as it enables them, in almost a single movement, to steal data from a vast swathe of victims, increasing the odds of a payout. In the past WithSecure have observed targeting of numerous other vulnerable file transfer services, including WS_FTP, GoAnywhere MFT, Citrix File Share, and Accellion FTA.

3.3 Vanishing Github comments allow for stealthy malware distribution

Researchers discovered a campaign that abused “ghost” Github comments on legitimate repos to host malicious software. The researchers found that it is possible to attach files to comments on Github repos, even if those comments are not actually posted. When you add a file to a draft comment it is uploaded to the Github CDN and is accessible via a URL under the parent repo. However, this occurs when the file is attached to the comment, before the comment is actually posted. As such, it is possible to add files to a draft comment, then discard the comment. The files will still exist in the GitHub CDN even after the comment has been deleted, and while accessible from their specific upload URL, they will not be linked to from anywhere else. There is no way to delete these files from your repo, and no way to prevent this type of upload, unless you block all comments on your repo, something which Github only allows for 6 months at a time on public repos. As software supply chain attacks are becoming increasingly common it is likely there is an increased threat to services such as Github and Gitlab.

WithSecure Insight

Attackers and defenders are in a constant arms race of stealth and detection. The ability to abuse legitimate services and infrastructure to host malicious payloads is constantly sought after and abused by attackers. While the use of this method is presently rare, it indicates the resourcefulness of attackers in their search for stealthy infrastructure to use in their attacks.

3.4 Espionage actors compromised Cisco ASAs with zero-days

This month it was disclosed that an espionage campaign that could not be linked to any previously known threat actors has been exploiting zero-days in Cisco ASA firewalls since as early as July 2023, up until early 2024. The actor used the compromised firewalls for initial access, reconnaissance, and traffic capture and exfiltration. The actor seemed to have an interest in Microsoft Exchange servers and network infrastructure devices from multiple vendors. This campaign was identified after Cisco were notified of security concerns by one of their customers in early 2024, and after a multi-month investigation involved several organizations that Cisco describes as intelligence partners. Cisco issued patches for two zero-day vulnerabilities in their ASA firewalls CVE-2024-20353 and CVE-2024-20359, in late March and early April, though they were unable to identify the initial attack vector.

WithSecure Insight

Cisco dominates the market for enterprise networking hardware, so to find out that a zero-day in their Firewall products was under active exploitation for over 6 months is extremely concerning. As repeatedly stated previously, infrastructure devices such as networking hardware may well be forgotten about as long as they keep functioning, and as a result may not be regularly patched. This story shows that patching these edge infrastructure devices is really a critical security requirement.

3.5 ChatGPT-4.5 can create CVE exploits from advisories

In a fascinating yet disturbing piece of research, it has been found that ChatGPT 4.5 can autonomously exploit vulnerabilities in real world systems with an 87% success rate if given a CVE advisory that describes the flaw. While other LLMs were tested, none were successful, although the researchers note that they did not have access to the two main ChatGPT competitors, Anthropic’s Claude 3 and Google’s Gemini 1.5. The LLM agent was created by using a LangChain ReAct automation framework/agent. The researchers noted that when the agent did not have access to the CVE description the success rate fell from 87% to 7%, though it’s not clear exactly what information the agent had in that case.

WithSecure Insight

The race between those seeking to exploit vulnerabilities and those seeking to patch vulnerabilities is increasingly being won by the wrong team, and this will not help.

3.6 Threat actor implements LLM generated dropper

Researchers at ProofPoint believe they have identified a Powershell dropper used by threat actor TA547, an Initial Access Broker (IAB), which was generated using an LLM. The dropper was delivered via phishing emails which contained password protected ZIP files, the ZIP files contained LNK files, and the LNK files executed the Powershell script to drop the Rhadamanthys infostealer. The LLM generated dropper wasn't identified due to any amazing, inhuman capabilities or obfuscation, but rather by the fact that every single component of the script was preceded by a hashed comment which very specifically described the function and intent of every component piece of code. This is not commonly seen in malware droppers but is very typical of LLM-generated code.

WithSecure Insight

It appears that as yet the AI-pocalypse has not yet come. This is one of the few indications we have seen that threat actors have been using LLMs for malicious purposes in real world attacks. It is interesting that the dropper was not particularly special or unique, apart from the verbose comment format. How possible is it to detect LLM generated attacks beyond something as simple as this? If targeting, lure generation, or data processing is being performed with LLM agents, how obvious would that actually be at the sharp end of an attack campaign? It may well be that the benefit of AI to cyber criminals will be in increasing efficiency gains, not in some AI cyber weapon of mass destruction.

3.7 Mass malicious phishing service taken down by law enforcement

The phishing as a service platform LabHost was taken down by law enforcement in a global operation which compromised the infrastructure and led to 37 arrests so far, including the original developer.

LabHost charged cyber criminals a monthly subscription to access the service, which provided access to a variety of different phishing kits which could specifically target North American banks and services with phishing pages and automatically generated and distributed phishing emails. The service gave even low skill actors access to effective phishing tools, including Attacker in the middle capability to capture MFA tokens. The takedown and associated investigation of LabHost identified 40,000 phishing domains which were being used by the roughly 10,000 users of the site. Users paid an average of $250 per month, and at the time of the take down the operators of LabHost are estimated to have received over $1 million in subscriptions. Investigators state that LabHost was used to steal approximately 480,000 credit cards, 64,000 PINs, and 1 million passwords.

WithSecure Insight

It is excellent to see another law enforcement takedown operation take effect. While phishing as a service may not have been reported in the same way as ransomware as a service, it is a significant part of the modern cyber-crime industry. The numbers given here really illustrate the impact of this service, which offered really quite advanced phishing capabilities to any actor with the money to pay for them.

3.8 Juniper decides to patch Junos

This month Juniper patched Junos, their Linux based operating system that runs on Juniper Cloud Native Routers and in Juniper cRPD (essentially a Junos Docker image), resolving vulnerabilities that are native to Junos, and also those from external software packages included in the OS. What is significant about this is that the update addresses 82 separate CVEs. Some of the lower severity vulnerabilities that were patched date back to 2011, which suggests that possibly they have not updated the software since then. Among these vulnerabilities were six 9.8 severity CVEs which date back as far as 2019. There was also one new Junos native vulnerability that was patched, CVE-2024-30407. This was the use of a hard coded private key in Junos which would allow AiTM attacks to undetectably intercept SSH traffic, resulting in complete compromise of the device.

WithSecure Insight

While exploitation of this vulnerability would be situational, requiring an attacker to have achieved a certain network positioning already, that is a truly fundamental security issue.

3.9 D-Link NAS devices have a hardcoded backdoor

A vulnerability has been discovered which affects multiple end-of-life D-Link NAS devices. The issue is a hardcoded account with an empty password which is present on these devices, and a command injection vulnerability in a CGI script. Together these issues enable unauthenticated RCE. While these devices are end of life products, 92,000 of them are still Internet accessible. D-Link have stated that they will not be addressing the vulnerability as the devices are not only End of Service, but End of Life, and have been for some time. As such the manufacturers recommendation is to simply retire them and replace with newer devices.

WithSecure Insight

While this is sensible advice, and it's hard to see what else D-Link could do about the issue, it does mean that there will be another 90,000 or so vulnerable devices connected to the Internet, waiting to get co-opted into a botnet.

3.10 JetBrains fixes 26 security problems

We reported on the JetBrains/Rapid7 disagreement last month around patching transparency. JetBrains appear to have made their stance on the matter quite clear this month by issuing a patch advisory that states that "26 security problems have been fixed", with no further details.

WithSecure Insight

While JetBrains head of security has said that the language used is because "most" of the problems relate to upstream libraries, and are not relevant or exploitable in TeamCity, one would hope that there might still be a middle place to be found somewhere between this and Rapid7's total transparency.

3.11 PyPi forced to disable new sign ups and project creations due to volume of malicious activity

A single threat actor's campaign of automated sign-ups and malicious new project creations in order to typo-squat benign projects reached such a volume on the PyPi Python package repository that they were forced to actually disable all sign ups and project creations for roughly 10 hours. Typo squatting means to register a very similar name to a legitimate resource, in the hope that people trying to access the legitimate resource will typo the name and end up at your malicious resource.

WithSecure Insight

WithSecure have noted an increase in attacks targeting and using software package repositories in 2024. This drastic and public response by PyPi illustrates just how big an issue this activity has become for them.

3.12 Sisense customer data stolen, more than a thousand major organizations affected

Sisense is a business intelligence company that offers a "fusion intelligence platform". This platform allows companies to link multiple third-party business intelligence services to the Sisense platform and view and combine them in a single location. This month CISA issued an advisory that Sisense had suffered a breach, and strongly recommended that all Sisense customers reset any and all credentials or secrets that had been shared with the company. Sisense has over 1,000 customers, many of which are large enterprises. Reports suggest that the attackers accessed Sisense's self-hosted GitLab where they found a credential that gave access to Sisense's Amazon S3 storage buckets. They then exfiltrated multiple terabytes of Sisense customer information, which almost certainly included large numbers of access tokens and credentials which the customers would have supplied to Sisense to access and process their data.

WithSecure Insight

The full impact of this intrusion may not yet be realized due to the theft of authentication material. It is another example of supply chain exposure, a risk that is extremely difficult for a CISO to terminate.

4 Threat data highlights

4.1 Exploits

Looking at the highest increases in WithSecure exploit detections this month, there was a huge surge in a historic VBS exploit (CVE-2006-3649), increasing from 1 detection to 5,500. There was also a 10x increase in exploitation of the 2023 WinRAR RCE (CVE-2023-38831). Multiple modest volume MS Office vulnerabilities saw increases of 20-50%, while a 2021 MS Exchange RCE (CVE-2021-27065) tripled in volume from 53 to 154.

Looking at the largest decreases in WithSecure detections, we see a similarly huge drop in detections of a Java 6 RCE (CVE-2010-0840). Two other java vulnerabilities (CVE-2010-4452 and CVE-2008-5353) have dropped significantly, though in smaller volumes. As well as this, all of the PDF related vulnerabilities from last month have dropped to 0 detections this month, which does reinforce the idea that they were all caused by a set of similar files.

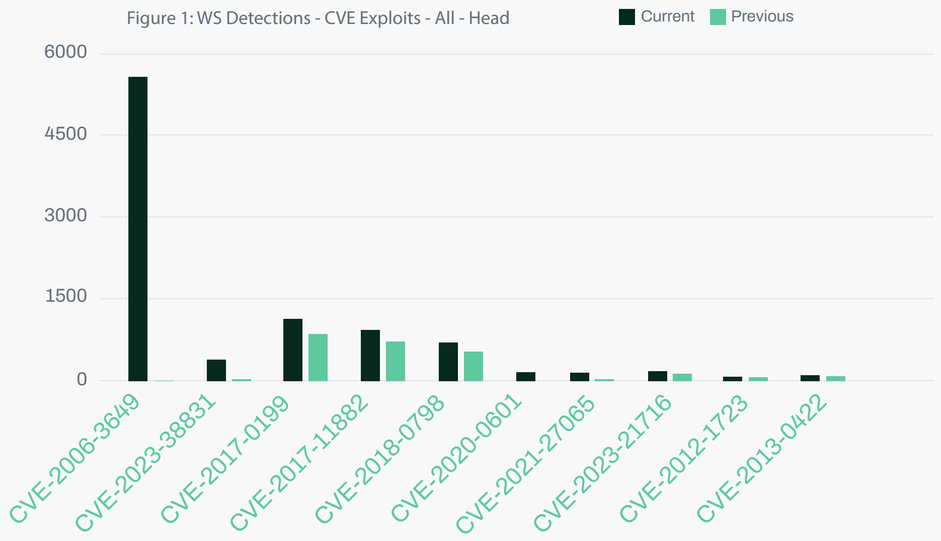

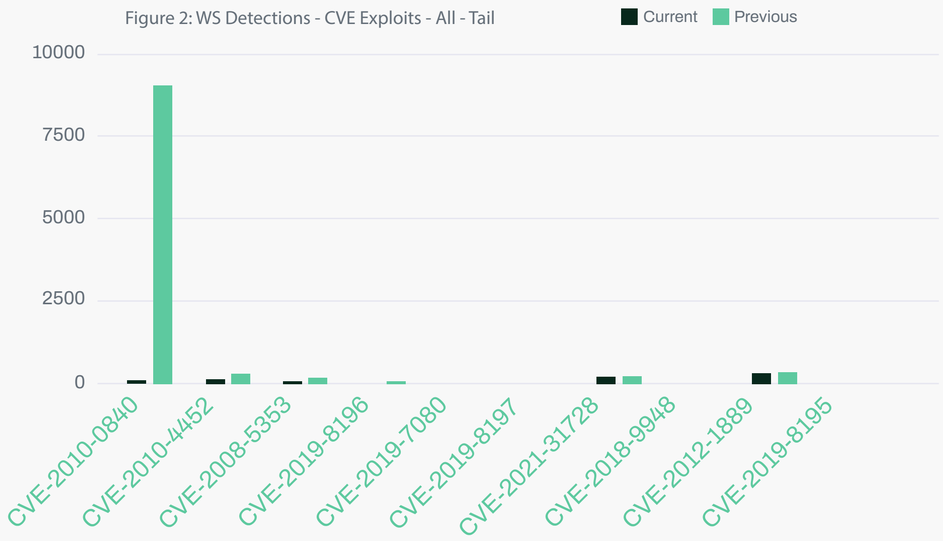

Looking at VirusTotal detections of 2023/24 CVEs with the largest increases this month there was a 16x increase in the number of files detected with the XZ Utils backdoor. Two 2023 Windows Privilege escalation exploits (CVE-2023-21768 and CVE-2023-28252) saw almost 5x increases in volume compared to the previous 30 days, while Microsoft's new Proxy Driver Spoofing Vulnerability (CVE-2024-26234) saw 12 detections this month. That is most likely based on detections of the legitimate code signing drivercertificate that was used to sign malicious files. There have also been detections of the new Palo Alto Panos GlobalProtect unauthenticated privileged RCE (CVE-2024-3400).

Looking at VirusTotal data for detection drops in 2023/24 CVEs, there were significant drops of around 50% in a Windows MSHTML privilege escalation vulnerability (CVE-2023-32046), and one of the Windows .URL file SmartScreen bypass vulnerabilities (CVE-2023-36025). While they have dropped hugely, the detection volumes remain high, at roughly 2,500 for each vulnerability. This mirroring between the two does raise the question as to whether the two are related somehow, for example being commonly used together, but that would only be speculation based on this data. The other, 2024 Windows .URL SmartScreen bypass vulnerability which was added to the KEV in February has also dropped significantly this month (CVE-2024-21412), dropping from 941 detections to only 240.

Figure 3: VirusTotal 2023/24 CVE Exploit Detections - Head

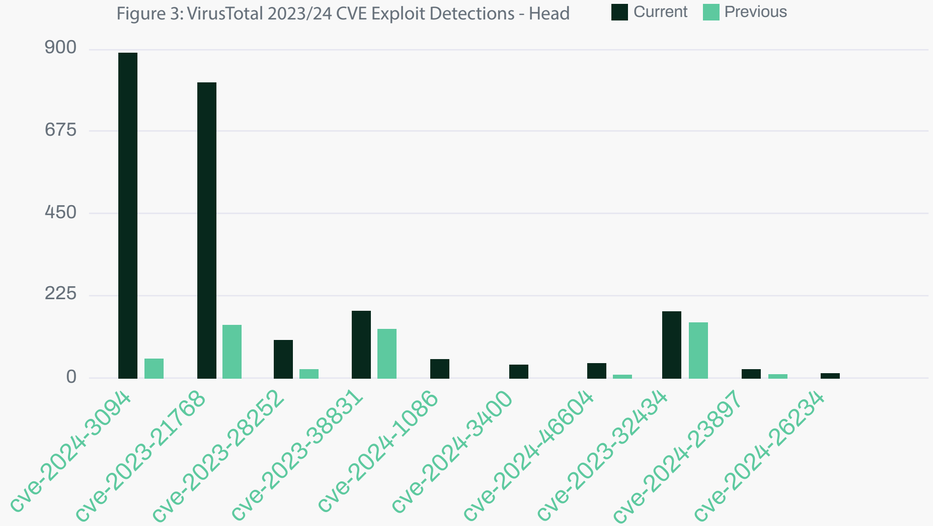

Figure 4: VirusTotal 2023/24 CVE Exploit Detections - Tail

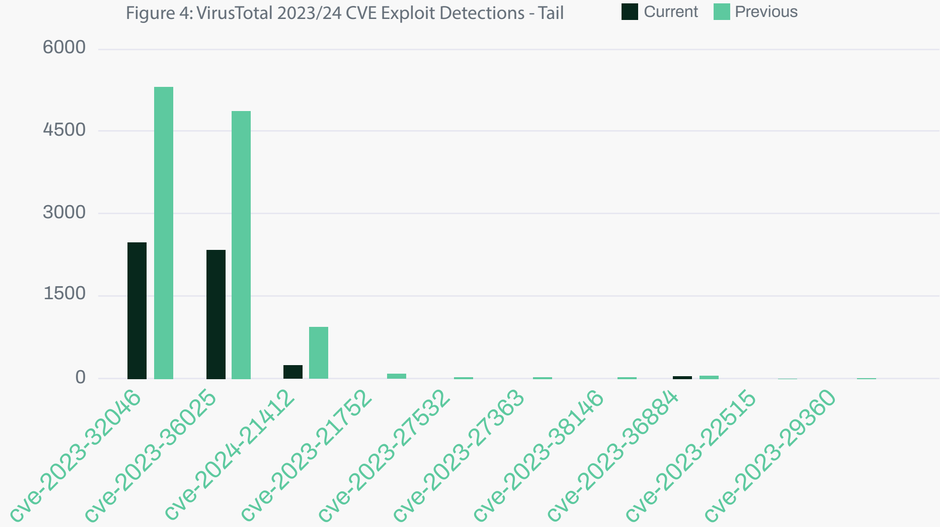

In all time CVE vulnerability detections from VirusTotal there has been quite an interesting development. CVE-2017-17215 is a Huawei HG532 router unauthenticated RCE which made it onto the lowest rung of this graph last month when detections doubled in volume from ~4,000 to ~8,000. In this month's data it is the most detected and fastest growing CVE by far, having increased from ~8,000 to ~16,000. This suggests that there may be a new and sustained campaign targeting this vulnerability.

Digging into the data about the Huawei exploit detections, we can look at submissions from the last 60 days and extract the submission location, although this location data is only available for around 30% of submissions. Do note that this is a stacked line graph, just to make a few things slightly clearer:

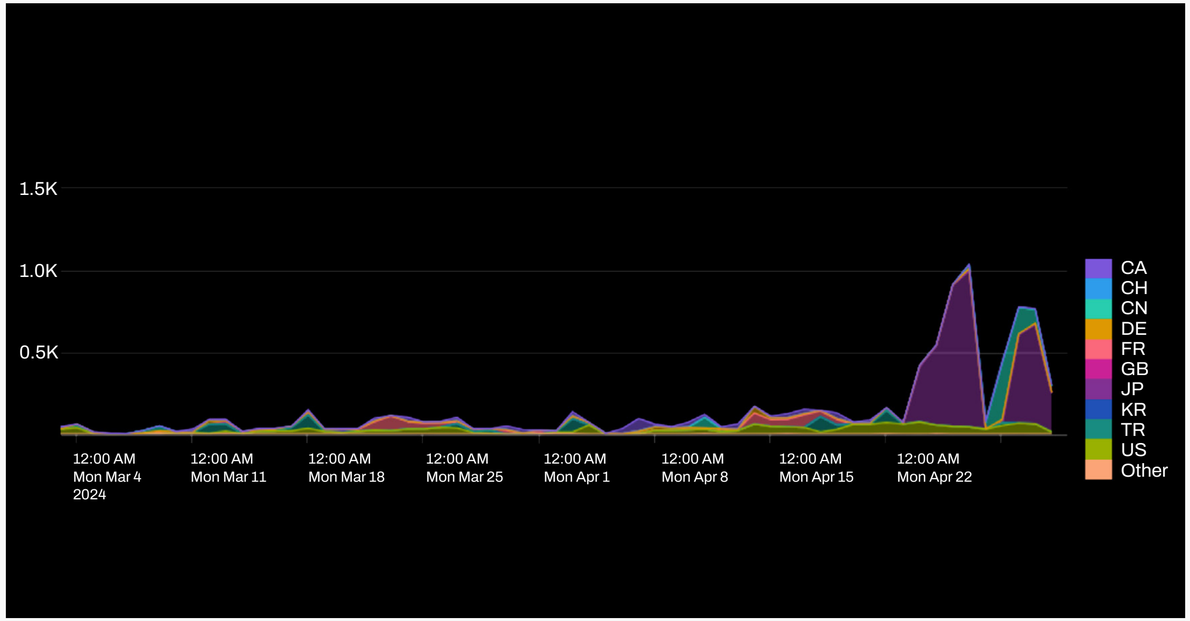

This shows a clear spike in submissions from Japan in late April, but also shows that over this time there were submissions from China, the US, Turkey, and multiple EU countries:

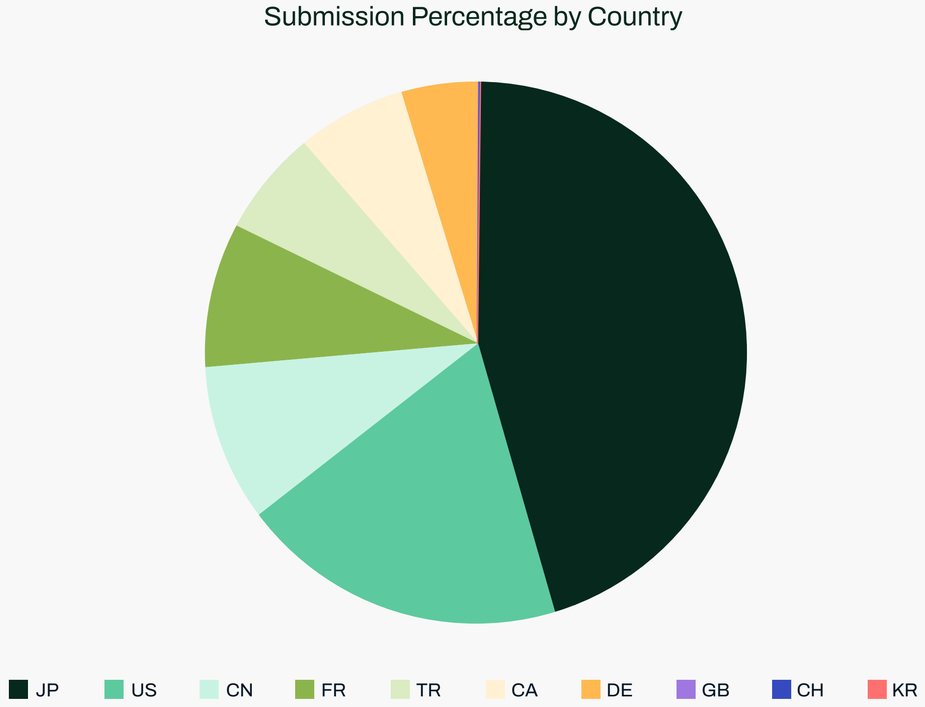

This is quite an interesting range of countries which at first glance does not seem to align with any single geopolitical event, though of course financially motivated attackers will also attempt to compromise routers for use in botnets.

Finally, looking at VirusTotal exploit detection drops for all CVEs shows interesting activity slightly further down the graph. A windows 10 privilege escalation vulnerability (CVE-2021-1732) has dropped from over 12,000 to 763 this month. This vulnerability spiked to that high number of detections last month but has now returned to normal levels again, which suggests it was being used in a short-lived mass campaign of some kind.

The Apache Tomcat HTTP/2 DoS vulnerability (CVE-2019-0199) which we observed increasing last month has decreased by half this month, which suggests that may also have been a short-lived exploitation campaign. There are also two encryption protocol vulnerabilities (CVE-2013-2566 and CVE-2012-5691) in the stats this month, one relating to IKE and one to use of RC4 in TLS/SSL. These have identical numbers of detections, so it is very likely that they are being triggered by the same files. It is however interesting to see detections of such old and specific vulnerabilities. While they are dropping in comparison to last month, there were still 3,000 detections this month, and almost 5,000 last month.

Figure 5: VirusTotal CVE Exploit Detections - Head

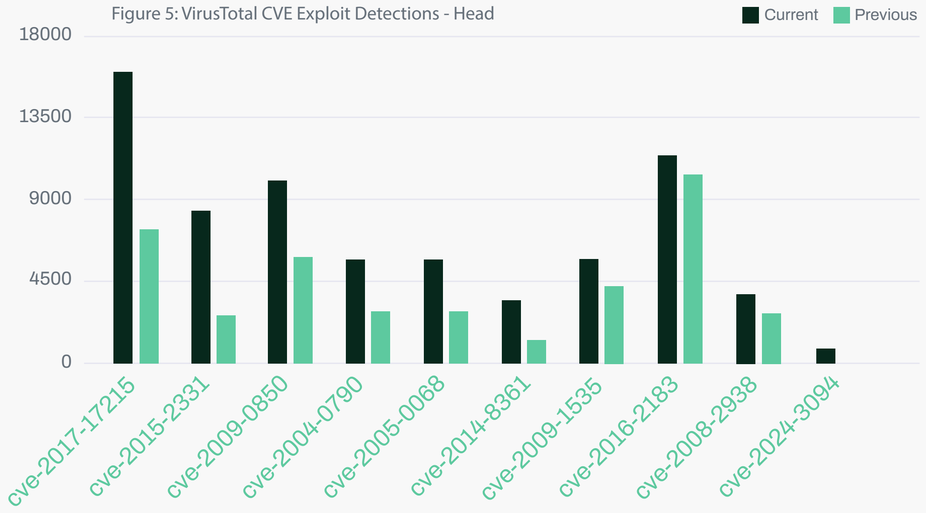

Figure 6: VirusTotal CVE Exploit Detections - Tail

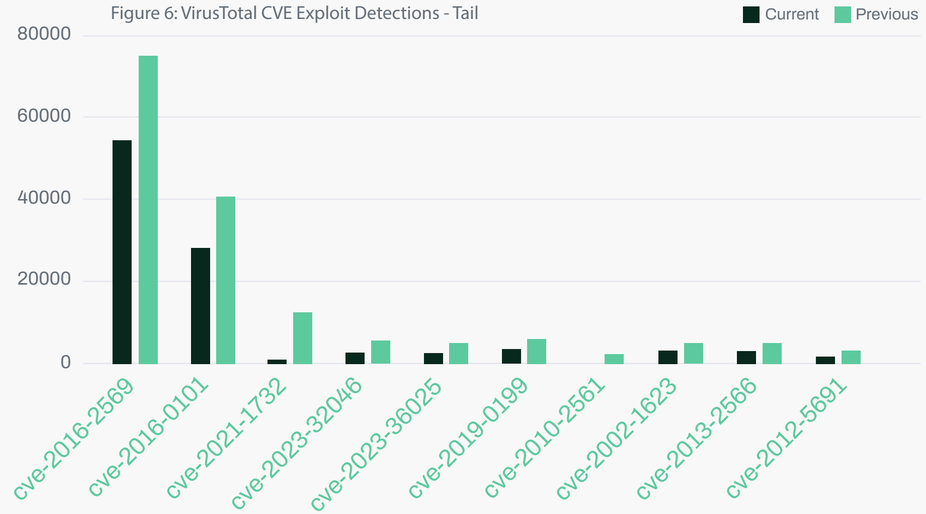

4.2 Newly exploited vulnerabilities

Looking at the additions to the KEV this month, we can see the CrushFTP and Cisco ASA zero-days that we discuss this month.

CVE ID	Vendor	Product	Vulnerability	Date added	Description
CVE-2024-29748	Android	Pixel	Android Pixel Privilege Escalation Vulnerability	04/04/2024	Android Pixel contains a privilege escalation vulnerability that allows an attacker to interrupt a factory reset triggered by a device admin app.
CVE-2024-29745	Android	Pixel	Android Pixel Information Disclosure Vulnerability	04/04/2024	Android Pixel contains an information disclosure vulnerability in the fastboot firmware used to support unlocking, flashing, and locking affected devices.
CVE-2024-3273	D-Link	Multiple NAS Devices	D-Link Multiple NAS Devices Command Injection Vulnerability	11/04/2024	D-Link DNS-320L, DNS-325, DNS-327L, and DNS-340L contain a command injection vulnerability. When combined with CVE-2024-3272, this can lead to remote, unauthorized code execution.
CVE-2024-3272	D-Link	Multiple NAS Devices	D-Link Multiple NAS Devices Use of Hard-Coded Credentials Vulnerability	11/04/2024	D-Link DNS-320L, DNS-325, DNS-327L, and DNS-340L contains a hard-coded credential that allows an attacker to conduct authenticated command injection, leading to remote, unauthorized code execution.
CVE-2024-3400	Palo Alto Networks	PAN-OS	Palo Alto Networks PAN-OS Command Injection Vulnerability	12/04/2024	Palo Alto Networks PAN-OS GlobalProtect feature contains a command injection vulnerability that allows an unauthenticated attacker to execute commands with root privileges on the firewall.
CVE-2022-38028	Microsoft	Windows	Microsoft Windows Print Spooler Privilege Escalation Vulnerability	23/04/2024	Microsoft Windows Print Spooler service contains a privilege escalation vulnerability. An attacker may modify a JavaScript constraints file and execute it with SYSTEM-level permissions.
CVE-2024-4040	CrushFTP	CrushFTP	CrushFTP VFS Sandbox Escape Vulnerability	24/04/2024	CrushFTP contains an unspecified sandbox escape vulnerability that allows a remote attacker to escape the CrushFTP virtual file system (VFS).
CVE-2024-20359	Cisco	Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD)	Cisco ASA and FTD Privilege Escalation Vulnerability	24/04/2024	Cisco Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD) contain a privilege escalation vulnerability that can allow local privilege escalation from Administrator to root.
CVE-2024-20353	Cisco	Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD)	Cisco ASA and FTD Denial of Service Vulnerability	24/04/2024	Cisco Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD) contain an infinite loop vulnerability that can lead to remote denial of service condition.
CVE-2024-29988	Microsoft	SmartScreen Prompt	Microsoft SmartScreen Prompt Security Feature Bypass Vulnerability	30/04/2024	Microsoft SmartScreen Prompt contains a security feature bypass vulnerability that allows an attacker to bypass the Mark of the Web (MotW) feature. This vulnerability can be chained with CVE-2023-38831 and CVE-2024-21412 to execute a malicious file.

5 Research highlights

5.1 Kapeka: A novel backdoor spotted in Eastern Europe

An in-depth technical analysis of the capabilities of Kapeka, and its connections to the Sandworm group intended to raise awareness of the threat amongst businesses and governments. The report was released alongside associated artifacts and tools, including configuration extractors, and a network communication emulator and decryptor. The report comes as Mandiant 'graduate' Sandworm to APT44 and CERT-UA release a report on 20 implants impacting the country.

5.2 Abusing search permissions on Docker directories for privilege escalation

During a recent engagement WithSecure came across an unfamiliar configuration pertaining to /var/lib/ docker permissions. This, combined with a number of other lower risk issues resulted in an attack path that allowed privilege escalation to root from a low-privilege user. As this was non-standard and uncommon, the observations and methods for leveraging this particular weakness are being shared for the benefit of the wider community.

5.3 Domain-specific prompt injection detection

This article demonstrates a practical approach to detect prompt injection attempts in LLM applications using a domain-specific dataset. We fine tune a DistilBERT model to train a classifier that is able to differentiate between legitimate inputs and potential injection attempts.

Who We Are

WithSecure™, formerly F-Secure Business, is cyber security's reliable partner. IT service providers, MSSPs and businesses – along with the largest financial institutions, manufacturers, and thousands of the world's most advanced communications and technology providers – trust us for outcome-based cyber security that protects and enables their operations. Our AI-driven protection secures endpoints and cloud collaboration, and our intelligent detection and response are powered by experts who identify business risks by proactively hunting for threats and confronting live attacks. Our consultants partner with enterprises and tech challengers to build resilience through evidence-based security advice. With more than 30 years of experience in building technology that meets business objectives, we've built our portfolio to grow with our partners through flexible commercial models.

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

Our latest reports: Threat-Research