Threat Highlight Report

June 2024

WITH secure

Contents

Foreword

There have been significant incidents and events in a number of fields this month. In an interesting geopolitical and commercial upset, Kaspersky have been banned from the United States, which is going to cause some significant disruption globally, not only in the US. A new MOVEit vulnerability was disclosed and while initially believed to be relatively low severity, within hours it was found to be chainable with a zero-day in a MOVEit software dependency. This pushed the severity up to critical and raising the specter of a wider software supply chain zero-day. In a campaign that just keeps going as more and more incidents come to light, around 150 customers of the Snowflake cloud data analytics service have been targeted and compromised by attackers leveraging stolen credentials to steal terabytes of data and PII, with victims including several household names. There has also been new information published regarding Chinese targeting of Fortinet firewalls in 2023, which finds that they were exploiting a since patched vulnerability as a zero-day for months, and in that time they managed to compromise and install a rootkit on 14,000 devices. Compromises continued even after the patch was released, and there are now estimated to be 20,000 compromised devices which it is not possible to remediate. Our last major story is that US SaaS supplier to the car dealership industry CDK has suffered a BlackSuit ransomware attack which has paralyzed their customers. The significance of this is the impact, CDK customers make up 85% of that industry, or 1-2% of all US GDP. As such, for each week that passes with their services down there is a ~USD$8 billion impact. There are many more stories beyond the highlights this month, including enough AI cybersecurity stories to justify their own, well populated section once again.

- Stephen Robinson, Senior Threat Intelligence Analyst, WithSecure

1 Monthly highlights

1.1 Kaspersky software banned from US due to national security concerns

On Thursday 20th June, the US Department of Commerce Bureau of Industry and Security announced a ban on all Kaspersky affiliates, subsidiaries or parent companies from directly or indirectly providing anti-virus software and cyber-security products or services in the United States or to US persons due to the risk to national security. They also added the Russian Kaspersky corporate entities AO Kaspersky Lab and OOO Kaspersky Group, as well as the UK entity Kaspersky Labs Limited to a list of entities who have co-operated with the Russian military and Russian intelligence. Kaspersky will be allowed to continue to provide signature and codebase updates to AV products until 29th September 2024 at 00:00 EDT. While US customers can continue to use Kaspersky products, they must assume the full risk of doing so.

Kaspersky has offices in 31 countries, serving 400 million consumer customers and 270,000 corporate clients across 200 countries.

WithSecure Insight

Kaspersky has been a technically capable member of the global cybersecurity industry for a long time, but it has also been viewed with suspicion due to its association with the Russian government of Vladimir Putin. While Kaspersky was previously banned from US Government networks, this latest ban applies to business with all US individuals and entities. It seems unlikely that this is in response to recent activity or a single event. While Kaspersky are believed to have been involved in some fashion (intentionally or otherwise) with the leaking and abuse of NSA “Equation group” hacking tools, that was many years ago. It is possible however that this ban is a form of financial and industrial sanction against Russia. Much of the suspicion around Kaspersky seems to be because they are a financially successful Russian company which is central to the Russian high-tech economy, with cybersecurity and technical professionals circling between Kaspersky and other employers, including the Russian Government. As such, the intent of this ban could instead to reduce funds, knowledge, and experience from entering the country and benefiting Russia’s high-tech sector, rather than a direct and damning indictment of Kaspersky.

1.2 New MoveIT SFTP vulnerability under active exploitation within hours of being disclosed

On Tuesday 25th June at 10:59AM Progress Software disclosed the existence of CVE-2024-5806 in MOVEit Transfer. Soon after this a POC was published by researchers at WatchTowr. The WatchTowr POC chained the MOVEit Transfer vulnerability and an additional software supply chain vulnerability in IPWorks SSH that Progress appears to have been un-aware of at that time. Seemingly in response to this, Progress increased the severity of the CVE to 9.1 and stated that there was now a new unpatched risk "related to a 3rd party product". Shadow Server reported observing exploitation attempts against the combined MOVEit/IPWorks exploitation chain hours after the POC was published.

When chained, these vulnerabilities allow unauthenticated access to a MOVEit Transfer server. The IPWorks SSH library vulnerability alone can be used to harvest the NTLM hash of the server.

A patch is available from MOVEit transfer for CVE-2024-5806, and MOVEit was contacting customers directly urging them to patch for some time before the disclosure. However, the vulnerability in IPWorks SSH library was a zero-day which still does not have a CVE.

On the 28th, three days after the initial disclosure, Progress updated their MOVEit advisory again, stating that they had now confirmed with IPWorks that the MOVEit Transfer patch also mitigated the IPWorks SSH vulnerability.

WithSecure Insight

The Clop ransomware gang's MOVEit zero-day exploitation campaign in 2023 had a huge impact on significant organizations around the world, so there should be real concern amongst users of MOVEit Transfer and the cybersecurity industry regarding this new vulnerability. While Progress have been contacting customers directly and urging them to patch for some time prior to disclosure, it appears they were only aware of the flaw in their own code. They were not aware of the IPWorks SSH vulnerability.

When contacted by the media, IPWorks CEO stated that the vulnerability has been patched at some point since June 24th, but there has been no public communication about this, and no CVE has been registered. The second update to Progress's advisory states that the MOVEit Transfer patch mitigates the IPWorks SSH vulnerability, but this refers only to MOVEit Transfer's use of IPWorks SSH. There is still no information from n software about what versions of IPWorks SSH are or are not vulnerable to the forced authentication vulnerability.

An additional concern is that the IPWorks SSH library is used in other solutions, which turns this into a software supply chain issue. Anywhere that IPWorks SSH is in use as an SSH server which accepts SSH keypairs is very likely also vulnerable to the forced authentication NTLM hash harvesting vulnerability.

"/n Software", the developers of IPWorks SSH also supply other software products as part of the IPWorks brand/suite, and it is possible that these may also include the same SSH functionality and vulnerability, depending on how they are used.

1.3 Snowflake data warehousing customers targeted in data theft campaign

At least 165 customers of the data warehousing/Data analytics cloud provider Snowflake have been compromised, with many experiencing significant data thefts, with data volumes into the Terabytes. Public reporting of this activity beginning on 31st of May, and the compromises have been linked to financially motived threat actor UNC5537. There does not appear to have been a compromise of Snowflake themselves, but instead a targeted campaign against Snowflake customer accounts, leveraging credentials stolen through infostealer infections over the previous 4+ years. Data leaks so far which appear to be from Snowflake customer accounts include Santander, Ticketmaster, Advance Auto Parts, Lendingtree, and Cylance. Analysis by TechCrunch found that more than 500 credentials for Snowflake customer environments were found in infostealer logs for sale on the dark web, including credentials for some of the known compromised organizations. In response to the incident, Snowflake have stated that they are in the process of requiring all customers to implement MFA.

WithSecure Insight

While it is positive that Snowflake do not appear to have been compromised and turned into a vector of supply chain compromise, this incident appears to have been enabled by the fact that Snowflake did not require customer accounts to use MFA by default, and there was no way for a customer organization to require/enforce MFA on all individual accounts, individual users would have to go into the settings and turn it on for themselves. In the situation that an infostealer infects a user and steals their username and password, which is all that is needed to compromise an organization, unless some form of MFA is required.

1.4 In 2023 a Chinese APT compromised more than 20,000 Fortinet infrastructure devices in a wide ranging zero-day campaign

On 11th of June, The Dutch Military Intelligence and Security Service (MIVD) issued an advisory stating that a Chinese APT espionage campaign which they initially disclosed in February was actually far larger than previously known. CVE-2022-42475 in Fortinet FortiOS and FortiProxy devices was exploited as a zero-day for a number of months, and during that time alone 14,000 Fortinet devices were compromised and infected by the actor. Targets included dozens of Western governments, international organizations, and a large number of companies in the defense sector. In total, the MIVD discovered that at least 20,000 devices were compromised. In addition, the rootkit like malware that was deployed, named Coathanger, is able to persist through system reboots, firmware upgrades, and even security updates, essentially giving the actor permanent root access to those systems.

WithSecure Insight

This is yet another significant mass exploitation campaign targeting infrastructure devices, and it highlights just how severe this kind of compromise can be. The very locked down nature of such devices limits both the information and tooling that is available to defenders for detection and remediation of compromise. These Fortinet infrastructure devices were essentially rootkitted, and so it may not even be possible to tell whether a device has or has not been compromised. One thing that is for certain is that it appears that much like with the Barracuda Email Security Gateway incident in mid-2023, the only guaranteed resolution is to decommission, replace, and then destroy the devices. Of course, Fortinet are a well-respected, major network security industry player who a huge number of organizations rely on to provide security solutions. So, if their devices cannot be trusted, what devices can?

1.5 Ransomware, espionage, and statistics

On 18th June the US based SaaS supplier CDK and its customers had their operations halted by a BlackSuit ransomware attack. CDK provide a full automotive dealership (i.e. car, heavy goods, agricultural, etc.) SaaS service including CRM, financing, payroll, support, vehicle servicing, inventory and parts ordering, and back-office operations to US clients. Following the attack, both of CDK's data centers were shut down, affecting all internal and external services. Customers were also advised to sever their always-on VPN links to CDK data centers. As a result of this attack car dealerships reported that they had resorted to pen-and-paper administration methods as they attempted to keep doing business.

On June 19th while CDK were attempting to bring their systems online again, they suffered a second cyber-attack, which does seem to indicate that they were attempting to bring their systems online too quickly, without sufficient security safeguards and diligence.

While CDK attempted to bring systems online again and, it is rumored, attempted to negotiate with the attackers, CDK customers received scam calls claiming to be from CDK support agents and affiliates. It is unknown if these were from opportunistic third parties, or if they were from the BlackSuit hackers.

As of the end of June CDK's services are still not available.

WithSecure Insight

While this may seem like a niche story, US automotive dealerships account for 1-2 percent of American GDP with an estimated turnover of around USD $500 billion (half a trillion). This incident has almost halted the business of around 85% of US dealerships for over two weeks, all from a single ransomware incident. As such the impact of this ransomware attack could be as much as USD$8 billion per week on the US dealership industry.

In 2022 CDK was purchased by a private equity group, and in 2023 many functions of CDK were outsourced to a 3rd party supplier. In unverified online discussions it has been stated that this included technical support and infrastructure. It is important to note that this does not automatically mean that there has been any kind of bad service or malfeasance, but in any new support contract or technical service there is the possibility of growing pains and miscommunications which could easily impact the service being delivered, giving an opportunistic attacker the vulnerability they need to initiate a compromise.

2 Ransomware: Trends and notable reports

2.1 The numbers

June 2024 saw a surprising and positive change in ransom-ware numbers. A sharp decrease in Lockbit victim numbers from 174 (May) to 9 in June has seen total numbers fall to 319 victims, the lowest on record since February 2023. As a proportion of total numbers, there were also sharp reduc-tions in 8Base, INC Ransom and RansomHouse. Cactus and relative newcomers SpaceBears increased victim numbers by 11 and eight respectively. Play posted the most victims this month, a slight increase on last month, followed by RansomHub and Akira.

While the numbers of victims posted was the lowest in over a year, the number of unique ransomware breach sites posting victims was also the lowest it has been since 2023 with 40.

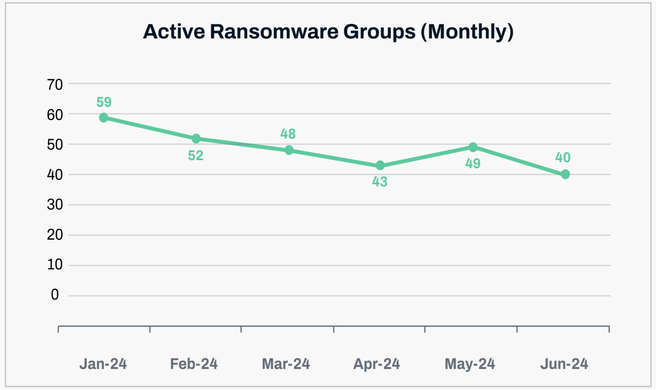

Ransomware	Count	Change
8BASE	8	-13
Abyss	4	3
Akira	20	-
Arcus Media	14	3
BianLian	9	-5
BlackBasta	15	-2
Blackbyte	1	1
Blacksuit	14	-1
Cactus	18	11
Cicada3301	4	4
CL0P	3	-
Cloak	6	3
Daixin	1	1
dAnon	3	-1
DarkVault	10	7
DragonForce	7	-6
Embargo	2	-1
Eraleignews	6	2
Everest	4	-1
FSOCIETY	2	-3
Hunters International	8	-3
INC Ransom	18	-15
LockBit	9	-165
Mallox	3	2
Medusa	18	-6
MetaEncryptor	1	-3
Money Message	1	-
Monti	3	-1
Play	35	3
Qilin	15	-4
Qiulong	1	-
Ransomhouse	3	-8
RansomHub	25	-2
Red Ransomware	1	-2
Rhysida	6	-
SenSayQ	2	2
Space Bears	10	8
Trinity	3	3
Underground	2	-1
WikiLeaksV2	4	4

2.1.1 Monitoring LockBit

Following the Law Enforcement action against Lockbit and the individuals working under the Lockbit banner, victim numbers posted to the Lockbit data leak site have been difficult to predict. WithSecure noted Lockbit's victims halving, and then halving again in the two months following the LEA action, but this month their numbers have been reduced to only nine new victims, far lower than Lockbit's numbers have been since they started.

There were a number of unusual changes to Lockbit infrastructure in June. Lockbit's telegram channel was reportedly seized by the FBI. This is unlikely to be the truth due to the broken English in the post explaining the seizure. Following this, there were some changes to Lockbit's DLS, and DLS mirrors where it appeared testing was underway in production on a new site with new DDoS protections. Following some test posts (victims named '12345.com' etc.), a victim was posted named 'US Federal Reserve', which turned out to be Evolve Bank & Trust. As there is no real reason to lie about what would be a perfectly viable victim, it has been speculated that the false posting came as a product of ignorance or carelessness.

We track Lockbit's victim numbers on a month-by-month basis, but it should be noted that that is simply an arbitrary period of time that means little to the threat actors. We have seen that the LEA action has impacted the Lockbit RaaS, but associated actors (particularly LockBitSupp) are still active and making changes to the Lockbit project. However Operation Cronos will affect Lockbit, it is currently clearly not in a 'business as usual' operational space.

Ransomware Data Not Deleted
WithSecure has long noted that organizations who have unfortunately been impacted by ransomware must consider the trustworthiness of the criminals they are dealing with. This month, the FBI revealed it had discovered definitively that Lockbit was holding stolen data it had claimed to have deleted on behalf of paying victims.

2.1.2 New groups

There were three new leak sites tracked this month: Cicada3301, SenSayQ and WikiLeaksV2*.

WikileaksV2
While it is a new data leak site, it does not appear that WikiLeaksV2 will serve specifically as a ransomware data leak site as many posts are duplicates of data leaked elsewhere. However, we have opted to include it as four of the 11 entries posted to WikileaksV2 have not been seen on other ransomware leak sites monitored. Of note, WikiLeaksV2 posted stolen data from Synnovis, the victim of a high profile Qilin ransomware attack which greatly impacted operations of several UK hospitals.

Cicada3301
Cicada3301, named after three famous internet puzzles and/or a group who sought to recruit those able to solve the puzzles, posted four victims, not enough to discern any patterns in victimology. It did post one UK victim, a financial software company with reported revenue of $1.4 billion USD.

SenSayQ
Only two victims were posted by SenSayQ, both small businesses with reportedly less than 200 employees. One from Italy and one from Lebanon.

2.2 Multiple London hospitals affected by Qilin ransomware compromise of supplier

Qilin ransomware actors seriously impacted an English NHS partner 'Synnovis labs' in a ransomware attack in June. Data relating to 300million individuals was reportedly stolen, and 400gb of data was posted to Qilin's ransomware leak site, their telegram channel and to a new data leak site called WikileaksV2. This comes after Qilin were almost certainly unable to successfully extort Synnovis for the demanded sum of $50million USD. Almost 200 cancer operations were postponed by the UK National Health Service. At one point blood testing was operating only at 10% of normal capacity.

Qilin started in late 2022 with its data leak site emerging in April 2023, and they have been posting steady, but low numbers to their leak site ever since. These numbers slightly rose following the Law Enforcement actions against Lockbit, and it has been noted that some affiliates may have turned from Lockbit to increase the work done with Qilin.

2.3 ScatteredSpider Linked to RansomHub

RansomHub is a relatively new extortion vehicle, its leak site starting in early 2024. It appears to be operated mainly from Russia. Symantec have noted that RansomHub is linked to an older ransomware variant, Knight. Ransomhub are aggressively attempting to lure affiliates from rival variants, particularly from now defunct brand ALPHV and Lockbit by allowing affiliates to directly collect a ransom and only then charging a 10% commission.

This tactic appears to have been successful with ScatteredSpider becoming a RansomHub affiliate, a prolific former Blackcat/ALPHV affiliate behind attacks impacting Okta, Twilio and Caesars Entertainment. ScatteredSpider themselves are more of a collection of cyber criminals working to multiple modus operandi, utilizing social engineering and publicly available tools and scripts to achieve their goals. In better news, it was also reported in June that an alleged leader of ScatteredSpider, a Scottish national who goes by the moniker 'tylerb' has been arrested in Spain.

2.4 TellYouThePass ransomware exploiting PHP vulnerability in new ransomware campaign

TellYouThePass ransomware has a lot of technical overlaps with HelloKitty Ransomware. In June, WithSecure research detailed problems with external service exploitation, particularly the increasingly short weaponization time for edge vulnerabilities often far outpacing patching cycles. This month it also became apparent that TellYouThePass ransomware has added an exploit to a PHP vulnerability to its arsenal. This is significant as there was only a ~48-hour gap between patch release, and exploitation by a prolific ransomware group – this was almost certainly galvanized and enabled by security researchers publishing proof of concept exploit code in line with the patch being published. This time period falls way inside even the most expedited organizational patching cycle.

3 Other notable highlights in brief

3.1 Microsoft and Google accused of GDPR violations by European Center for Digital Rights

The non-profit noyb (also known as the European Center for Digital Rights) has made a complaint to the Austrian Data Protection Authority to investigate Microsoft Education 365 for breaching the transparency provisions of GDPR, as well as collecting data on users regardless of ages. They state that neither Microsoft documentation, requests to Microsoft, or noyb's own research could clarify what data about children is being processed, and that Education 365 installs cookies without consent and uses them to analyze user behavior, collect browser data, and prepare advertising, regardless of the user's age.

This month noyb also made a complaint about Google's "Privacy Sandbox", stating that it is actually a tracking tool for targeted advertising that Google has presented as a privacy feature, which noyb believes does not meet Google's obligation for specific, informed and unambiguous consent under GDPR.

Both Microsoft and Google have made statements defending their products and actions.

WithSecure Insight

European privacy law, such as GDPR, is generally seen as stricter than US privacy laws. While EU companies may at times struggle to navigate the regulations, there is often a perception that the ethos of US companies regarding user privacy is more likely to come into conflict with GDPR. This may or may not be true, but it is interesting that two of the largest and most influential US tech companies, each of which sits at the center of the US (and international) tech industry ecosystem should both have complaints lodged against them in one month. It is significant that the complaints are lodged by noyb, as it was the lawsuits it brought that resulted in the EU Court of Justice striking down the "Privacy Shield" data protection arrangements between the EU and the US.

3.2 Microsoft delays debut of Recall after security and privacy concerns vociferously raised

Microsoft have announced that they will be delaying the rollout of the Recall feature which caused so much consternation last month. In addition, when it does rollout it will be as a preview available to users of the Windows Insider Program pipeline. It will no longer be enabled by default on all Copilot+ PCs.

WithSecure Insight

Many people will breathe a sigh of relief at this news, possibly including people within Microsoft. It was reported that Recall was developed so secretly inside Microsoft that the first that the external announcement was the first time many in Microsoft heard of it. It also seems relevant that Microsoft's decision on Recall came the day after Microsoft's President testified before US Congress that Microsoft's own security failings allowed Chinese threat actors to access US government emails, and also stated that Microsoft would implement every recommendation from the CSRB report into that incident. While we previously described Microsoft as a very large ship making a very slow turn, a better metaphor may be that it is a many large ships, all trying to turn at the same time.

3.3 Checkpoint VPN vulnerability under active exploitation

Exploitation of Checkpoint Security Gateway VPN CVE-2024-24919 began on May 31st, shortly after a POC became available. Attempts to target the vulnerability actually began a day earlier, coming from a Tiawan located IP address, but the exploit in those attempts did not work. By June the 5th ~782 IP addresses were scanning for and attempting to exploit the vulnerability, while at the same time over 13,000 Internet exposed Checkpoint firewalls were identified. 2% were identified by Censys as running a patched software version, while 4.6% were identified as running a potentially vulnerable version.

WithSecure Insight

This was yet another critical vulnerability in a firewall VPN gateway which could be exploited remotely, with no user interaction or privileges required. The one silver lining is that it only affected Checkpoint devices which were providing “Mobile Access” or IPSec-VPN services, however it is unclear just how much that reduced the potential exploitation surface. It is also worth noting that after last month’s advisory from the Norwegian NCSC to move from SSL-VPNs to IPSec-VPNs, this vulnerability does apply to IPSec VPNs.

3.4 Darkgate stops using AutoIt, moves to AutoHotkey

Darkgate is a modular MaaS infostealer and RAT which gained a lot of popularity with cyber criminals when its competitor, Qakbot, experienced interruption from a Law Enforcement takedown. Darkgate has used AutoIt scripts as part of the malware dropping process since its inception, but with the release of version 6, it now appears to have shifted entirely to AutoHotKey based scripts. At the same time a number of previous pieces of functionality have been removed, such as cryptomining, hVNC, and privilege escalation, possibly in an effort to become more difficult to detect.

WithSecure Insight

Cybersecurity is often described as an arms race between defenders and attackers, Darkgate is used by multiple attacker clusters, and delivery attempts are regularly detected and prevented. After almost one year it is interesting to see Darkgate move away from what has until now been a standard part of its infection process. It may be that the DarkGate developer felt the need to move away from AutoIt in order to improve the effectiveness, and reduce detections of DarkGate, but if so, this is simply part of the arms race.

3.5 Gitloker extortion campaign wipes GitHub repos

In an interesting twist on cyber extortion, an attacker is accessing GitHub repos with stolen/compromised account credentials, then deleting the contents and leaving behind a readme file stating that they have a backup and, of course, would be willing to supply it to the victim in return for payment.

WithSecure Insight

Cloud services are sometimes seen as ransomware proof, as there is a believe that any data uploaded to the cloud is always backed up somewhere. However, it is possible to intentionally delete data from GitHub without the possibility of retrieval. If sensitive data or secrets have accidentally been uploaded, this is of course a desirable feature, but it is interesting to see that ability being abused here. There are a number of advantages to this method for the attacker. There is no need for any custom software, simply standard Git commands. They also do not need to compromise a network, they simply need a valid GitLab account, which they could quite possibly find just from Infostealer logs. In that way, this is quite similar to the Snowflake incidents this month, and this type of targeting via 3rd party accounts could well be a trend to watch.

3.6 Comprehensive POC available for critical Veeam auth bypass

A POC for Veeam authentication bypass vulnerability CVE-2024-29849 was released on 10th June. The exploit works by sending a specially crafted VMWare Single Sign On token to the Veeam REST API on port 9398. The SSO token contains an authentication request (typically for an administrator) and an SSO service URL. The vulnerability is the fact that Veeam does not verify that the SSO URL provided is actually trusted in any way. The SSO token is simply sent to the specified SSO URL to verify that it is valid. As such an attacker can send a crafted token that points to a server they control, which will authenticate the SSO token and allow them to login.

WithSecure Insight

While the concept of Single Sign On is intended to increase security and accessibility, it does of course require correct implementation. It is an astonishing oversight to simply trust any SSO server that the client provides. So many vulnerabilities come about when a service is overly trusting of client/user provided input, yet the need to distrust and verify input is not some new revelation. The Perl taint checking functionality was first supported in 1989, 35 years ago.

3.7 BEC compromise leads to USD$445,000 loss

This month it was disclosed that the US town of Arlington, Massachusetts lost USD$445,000 to a business email compromise (BEC). Attackers had compromised town employee email accounts and were monitoring the emails being received. When the time was right, they created an email domain which impersonated a vendor working on a construction project and emailed the town authorities requesting a change in payment method to electronic transfer. As such the next 4 monthly payments were sent to the attacker's account, which only came to light when the vendor reported that they had not been paid for several months. Investigations found multiple additional fraud attempts which totaled USD$5 million, though none of the others were successful. The total cost of the construction project was USD$240 million.

WithSecure Insight

It is well understood that ransomware attacks can lead to huge monetary losses for victims, but while BEC is known to be a threat, it does not seem to get the same headlines, or be viewed with the same severity as ransomware. The possible financial losses due to BEC however are directly tied to the value of the communications that an organization engages in over email. Considering the quantity and value of business done via email then for an effective BEC attacker, the sky is the limit.

3.8 TeamViewer corporate network compromised, attributed to Russian APT

TeamViewer have stated that their corporate network was compromised by the Russian actor Midnight Blizzard/APT29. The information upon which that attribution is based has not yet been released. Because of the type of service TeamViewer offer, and the breadth of their userbase, there are of course supply chain concerns around the attack, however at present TeamViewer have strongly stated that they have no evidence that the TeamViewer product environment, or any customer data has been affected. They have stated that the compromise has been limited to the TeamViewer internal corporate network, which was accessed using a compromised/stolen employee credential. The attackers exfiltrated employee directory data including names, contact information, and encrypted employee passwords, which they state they have mitigated.

WithSecure Insight

It is good to see that once again, network segregation appears to have made the difference between compromise of an organizations corporate network or compromise of all of that organization's customers. However, the possible impact of a TeamViewer service/product compromise is such that many organizations will be on high alert after this news. The advice being offered to TeamViewer customers is really the standard best practice when using remote administration tools: require MFA, implement thorough logging and log monitoring, and enforce protections and restrictions around administrative access and activities.

3.9 Smishing attackers create DIY cell phone mast to flood nearby devices with malicious SMS messages

Two individuals have been arrested for setting up a DIY mobile/cell phone tower in London and using it to send thousands of scam SMS messages to any mobile phone that came within range. The arrests were made by the City of London police, and they describe it as a first of its kind crime in the UK. SMS phishing is often referred to as smishing, and typically involves sending messages across the mobile phone network from valid mobile numbers/SMS services. However, because mobile phone communication standards require mobile devices to authenticate to the network, but don't require the network to authenticate to the phone, it is possible to set up fake mobile phone base stations which can trick mobile devices into connecting to them and receiving SMS messages.

WithSecure Insight

While this may be a first of its kind crime in the UK, DIY phone masts (also known as IMSI catchers) have been a problem for some time in other countries, with criminal incidents in Vietnam, the Philippines, France, and Norway. SMS messages, whether sent through the legitimate mobile phone network or through fake base stations represent a security issue and targeting vector to beware of It is important to remember that SMS is a legacy piece of functionality, similar to email, but without the various security and authentication features such as DMARC. DKIM and SPF that have been added to enterprise email solutions over the years.

4 AI

4.1 Secrets accessed in compromise of Hugging Face Spaces

On May 31st, AI development platform Hugging Face announced that they had detected unauthorized access of secrets in their Spaces platform, and in response they revoked tokens which were present in the secrets that may have been compromised. This also pushed them to tighten up their security in general, as well as specifically for tokens, removing organization wide tokens and blanket read and write tokens in favor of more fine-grained access controls. They have also now implemented a KMS for secrets which they believe will improve their ability to identify and proactively revoke leaked tokens.

WithSecure Insight

Hugging Face is sometimes described as the GitHub of AI, where developers can access pre-trained models and “do AI”. As with any organization or platform that has had meteoric growth, there are very likely to be problems. One of the problems such rapid growth introduces is that what may have started out as a temporary solution to a problem for a small organization can rapidly turn into a major dependency that an entire service and software stack relies upon, increasing both the impact and the difficulty of applying a fix.

4.2 EmailGPT prompt injection flaw discovered

Prompt injection vulnerabilities have been identified in EmailGPT. EmailGPT is an API service and Google Chrome extension that used ChatGPT to assist users writing emails within Gmail. Researchers reached out to the developers of EmailGPT but did not receive a response within 90 days.

WithSecure Insight

This is quite a familiar story, Tool created which interacts with AI. Tool is then vulnerable to prompt injection. However the fact that this keeps happening is really the story here: Prompt injection is a known risk of using LLMs, and so anywhere that LLMs are being used, this risk needs to be guarded against.

4.3 Infostealer distributed via trojanized node/package for ComfyUI Stable Diffusion GUI

ComfyUI is a GUI front end for users of the Stable Diffusion AI image generator. Comfy AI allows users to construct image generation workflows by chaining together packages, or nodes. Developers can create and share custom nodes, and those nodes are just code that is executed by ComfyUI on the local workstation. It was discovered that a node named ComfyUI_LLMVISION was actually an infostealer trojan. A user on the ComfyUI sub-Reddit posted analysis of two different versions of the malicious code, each of which gathers credential information including API keys from the local system and sends it to a Discord webhook.

WithSecure Insight

The trust that is given when we download and execute somebody else’s code, whether open source or commercial, is often taken for granted. Often with mods and packages such as this there is some form of sandboxing to at least restrict the resources available to the downloaded code, but in this case it appears that ComfyUI nodes are simply able to run any code and local commands, and even make unrestricted external connections. When people say “AI is the wild west” this is the sort of thing they are talking about.

4.4 Sleepy Pickle, New ML model exploitation technique

Pickle is a format for saving Python objects to files and loading objects from files. A known issue with pickle files is that it is relatively easy for an attacker to insert malicious Python code/objects into the file. The Sleepy Pickle attack involves inserting malicious code into the pickle file of an AI model that will modify the AI model to behave in undesirable ways, such as inserting a backdoor, controlling the output of the model, or tampering with processed data.

WithSecure Insight

Sleepy Pickle is an interesting demonstration of the fact that quite a low-level attack on AI supply chains can have higher order effects on the behavior of models. The fact that Pickle files can be edited to contain malicious code is well known, but it is an impressive leap to go from that to subtly altering the behavior of an AI model. The most practical defense against this specific attack is not to use such an unsafe file format as Pickle, and instead to use a more secure format such as SafeTensors.

4.5 Meta to pause AI training with EU user data after privacy concerns raised by Irish Data Protection Commission

Meta announced that it will be delay training of LLMs on public content uploaded by Facebook and Instagram users in the European Union after receiving a request from the Irish Data Protection Commission (DPC). The issue appears to be that Meta had not sought explicit consent from users to use the data in this way, and explicit consent is a cornerstone of EU privacy law. Meta also received queries from the UK Information Commissioner's Office, as the UK currently operates under a similar privacy framework to the EU. The EU privacy organization that we mentioned earlier this month, noyb, has also filed a complaint against Meta alleging that they are breaking GDPR by making the data harvesting opt-out instead of opt-in, not explicitly stating what the data will be used for, and by extension not seeking explicit consent for those uses.

WithSecure Insight

Meta has described the complaints that have prevented the ingestion of EU data to train its AI models as shortsighted, preventing it from bringing certain features to the EU market. But at stake here is that users did not submit their data with the knowledge that it would then be used to develop and train AI models. Not only that, even knowing that data may be used to train AI models, there is currently no information from Meta as to what those AI models could be used for, so it would be very difficult for users to grant informed, explicit consent.

4.6 Microsoft publish details of Skeleton Key, a new prompt injection attack

Skeleton Key is a type of prompt injection attack detailed by Microsoft researchers which was found to work on AI models from Meta, Google, Anthropic, and OpenAI, except for GPT-4. It works by telling the model to modify its behavior when it encounters a previously defined guardrail, and instead of simply refusing to comply with the objectionable request, instead prefixing the reply with a warning message.

WithSecure Insight

Prompt engineering often seems like an exercise in creative logic, attempting to find situations and methodologies to either cause the model to disregard the instructions under which it operates, or to request the model to do something it is not supposed to do by complying with the letter of any restrictions, rather than the spirit. This type of vulnerability research does bear some resemblance to traditional unexpected input application hacking, which can often come down to how to bypass a developer specified blacklist, but it is a lot more like social engineering than programming.

5 Threat data highlights

5.1 Exploits

Very few all-time exploits saw a significant increase this month. The largest increase was for a 2008 Excel RCE which increased by around 200%. Interestingly however, Microsoft products made up only half of the vulnerabilities exploited here, with the other half made up of several Java RCEs, a Flash RCE and a Linux Kernel privesc.

Looking at falling detection numbers, there is a significant drop in a 2012 Windows Media Player crafted midi file vulnerability, and four Microsoft Office RCE vulnerabilities often used in malicious document phishing/spam. Two of the vulnerabilities were 2017 CVEs in Microsoft Equation editor, though the differing number of detections suggests they are not duplicate detections. There is also a proportionately significant drop in exploitation of the 2022 Apple OS privesc towards the middle of the graph.

Figure 1: WS Detections - CVE exploits - Largest increases

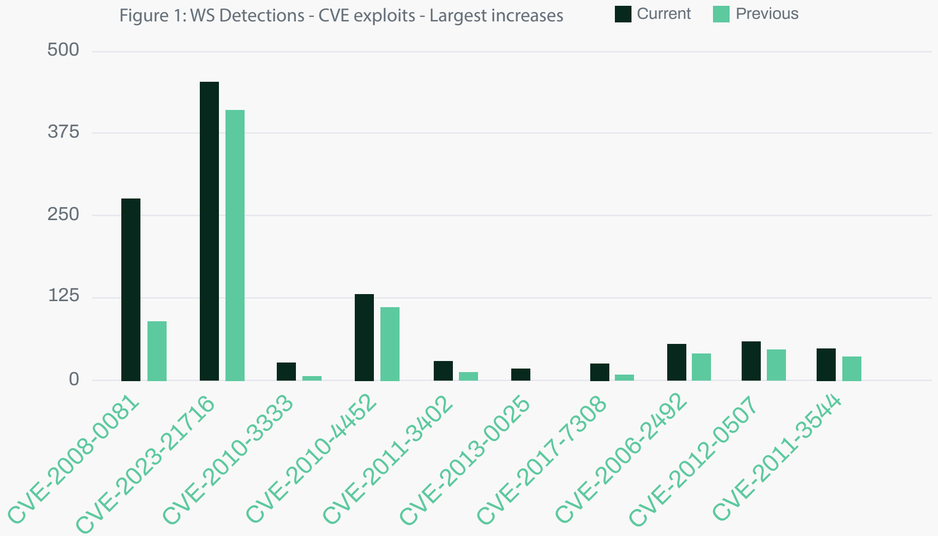

Figure 2: WS Detections - CVE Exploits - Largest decreases

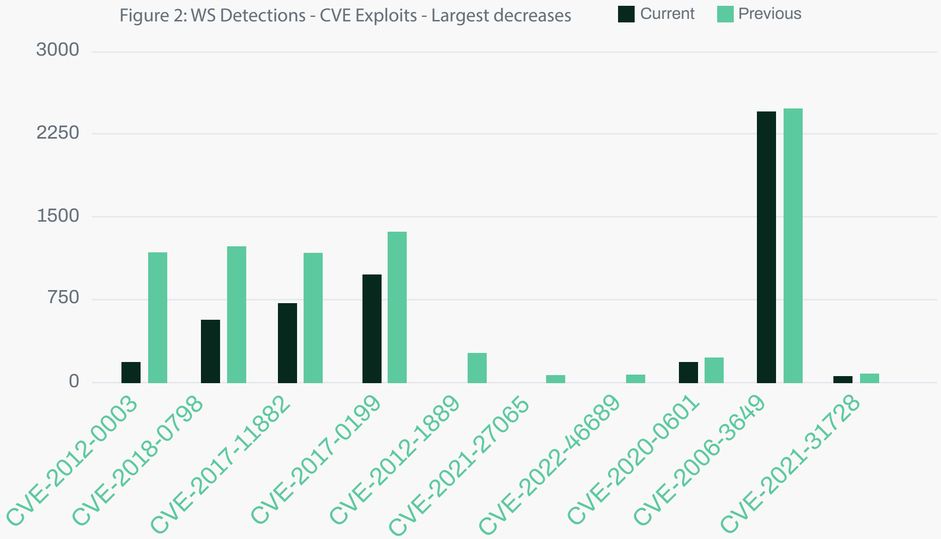

Looking at 2023/24 detections there were only three CVE exploits detected, each with relatively small fluctuations.

The volume of detections that were new for this month compared to last is very small, and most likely just random fluctuations. The only thing of interest is that 4 of these (including the highest entry) were Adobe Flash Player exploits.

Figure 3: WS Detections - 2023/24 CVE Exploits - Largest increases

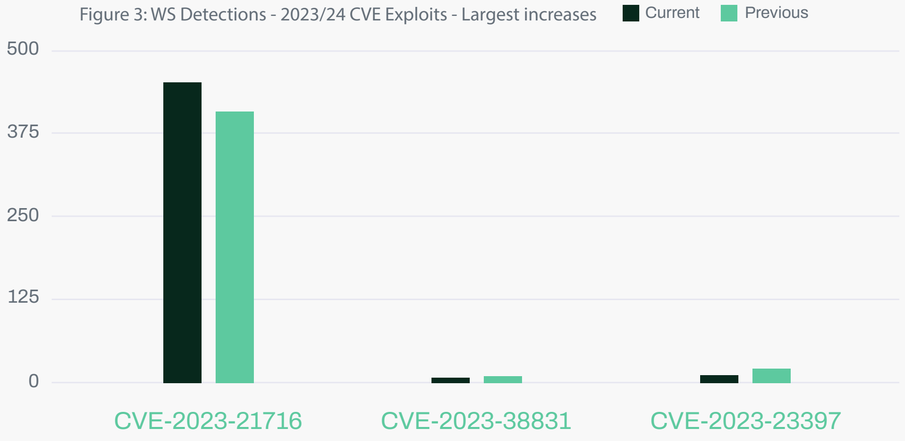

Figure 4: WS Detections - New in June vs May

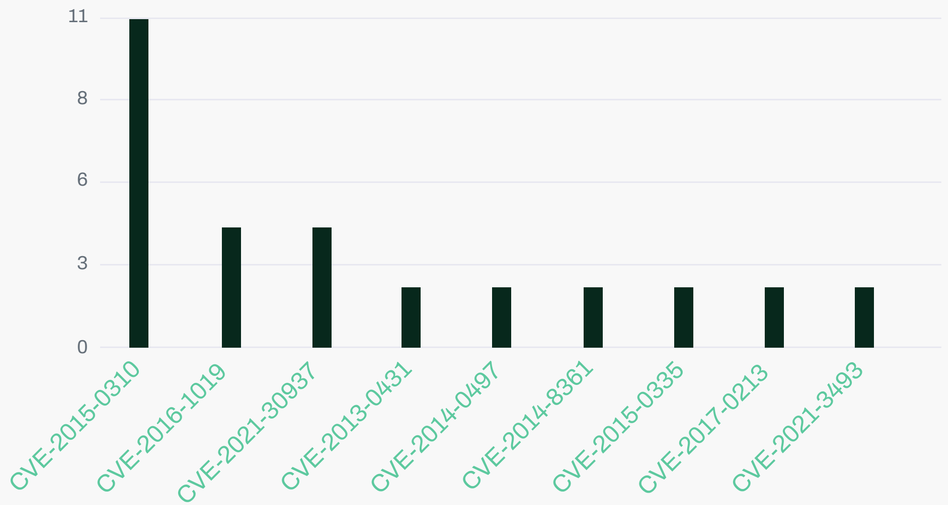

In VirusTotal data for increasing detections of 2023/24 exploits, the highest increase is one that was not seen last month, a 2024 Windows CSC Privesc. This vulnerability was patched and published in April, with POCs available in June. Interestingly at number 2 is an Ivanti ICS Auth bypass vulnerability, which indicates continued interest in Ivanti ICS from either researchers or attackers. A new Firefox RCE vulnerability, published in May, comes in at 4th place.

In decreasing VirusTotal 2023/24 detections are a number of Windows vulnerabilities, with the .URL file SmartScreen bypass and patch bypass coming in second and third. The XZ backdoor is in 4th place, once again most likely showing continued interest by researchers. In 8th place is the Palo Alto GlobalProtect unauthenticated root RCE, which has dropped almost to nothing. Considering the difficulty of extracting files from Palo Alto devices, it is interesting to see this in the statistics at all.

Figure 5: VirusTotal 2023/24 CVE Exploit Detections - Largest increases

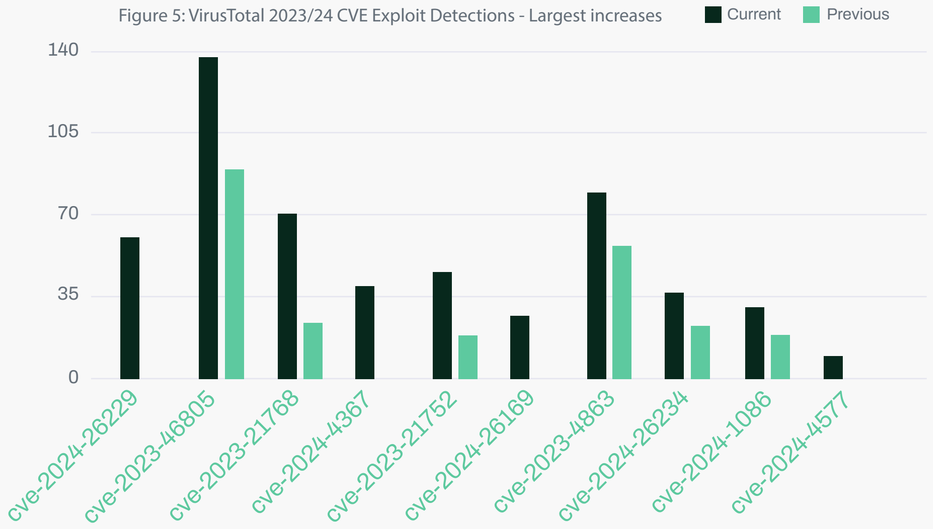

Figure 6: VirusTotal 2023/24 CVE Exploit Detections - Largest increases

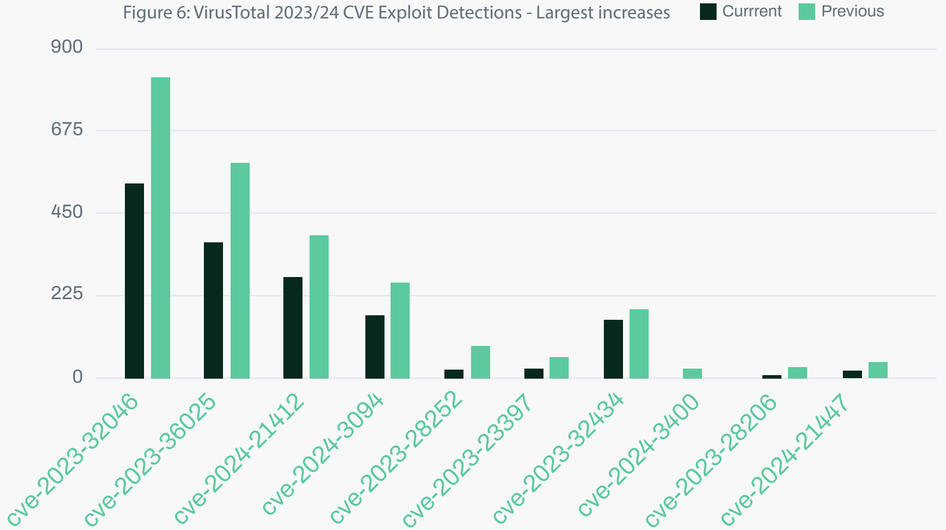

Looking at increases in VirusTotal detections for any CVE, there are huge increases detections of exploits for two very similar vulnerabilities, both DoS via crafted ICMP with Path MTUD, though at first place is that vulnerability for BIG-IP, and at second places is a multi-implementation/generic vulnerability. Then, at 3rd and 4th place come two more ICMP DoS vulnerabilities, but with identical detection numbers, indicating that they are most likely being triggered by the same files/code. Interestingly, that does mean that the top 4 detections here are all ICMP DoS exploits. In another interesting but easy to overlook data point, the final entry on the graph is a 1999 CVE remote auth bypass for wuarchive ftpd and ProFTPD, which has increased from 13 to 1009 detections.

Finally looking at VirusTotal all time CVE detection decreases, there are a number of large drops, with a Windows Media Player RCE at first place, though that still shows 12,000 detections this month. In second place however is a PHP libzip DoS and RCE exploit, which has fallen from 15,000 to 478. Interestingly, in the bottom half of the graph are 3 vulnerabilities with almost identical volumes which relate to DNS poisoning attacks against SQUID, BIND and Microsoft DNS servers. All of these have decreased from ~5,000 last month to ~1,000 this month.

Figure 7: VirusTotal CVE Exploit Detections - Largest increases

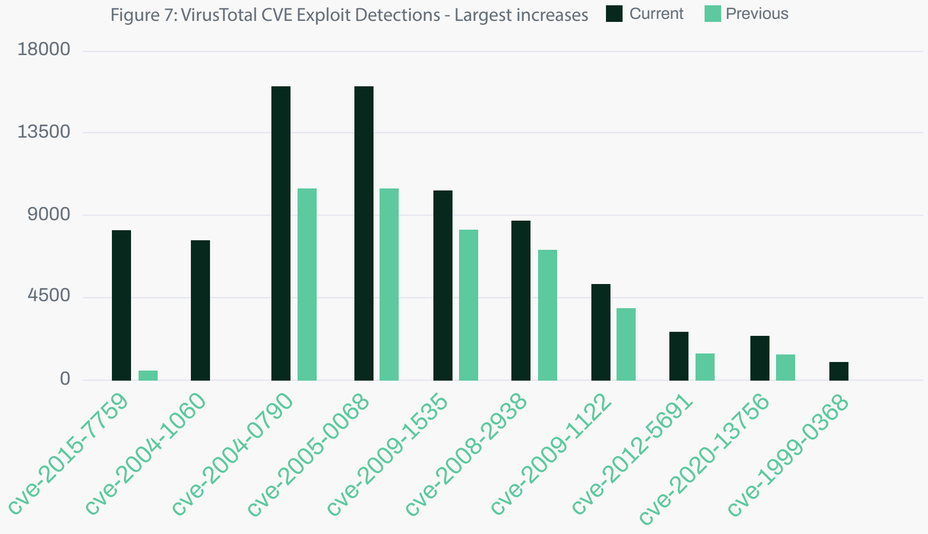

Figure 8: VirusTotal CVE Exploit Detections - Largest decreases

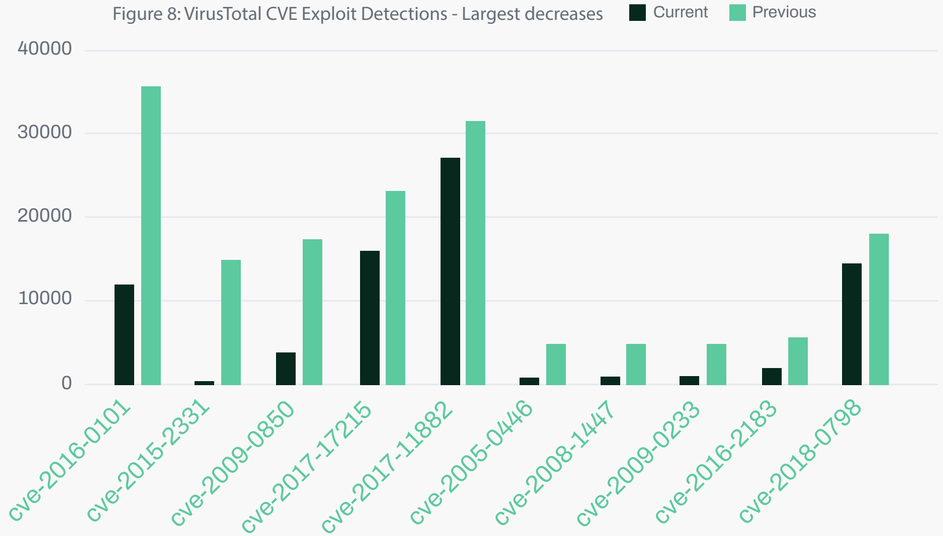

5.2 Newly exploited vulnerabilities

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-2017-3506	Oracle	WebLogic Server	Oracle WebLogic Server OS Command Injection Vulnerability	03/06/2024	Oracle WebLogic Server, a product within the Fusion Middleware suite, contains an OS command injection vulnerability that allows an attacker to execute arbitrary code via a specially crafted HTTP request that includes a malicious XML document.
CVE-2024-4577	PHP Group	PHP	PHP-CGI OS Command Injection Vulnerability	12/06/2024	PHP, specifically Windows-based PHP used in CGI mode, contains an OS command injection vulnerability that allows for arbitrary code execution. This vulnerability is a patch bypass for CVE-2012-1823.
CVE-2024-4610	Arm	Mali GPU Kernel Driver	Arm Mali GPU Kernel Driver Use-After-Free Vulnerability	12/06/2024	Arm Bifrost and Valhall GPU kernel drivers contain a use-after-free vulnerability that allows a local, non-privileged user to make improper GPU memory processing operations to gain access to already freed memory.
CVE-2024-4358	Progress	Telerik Report Server	Progress Telerik Report Server Authentication Bypass by Spoofing Vulnerability	13/06/2024	Progress Telerik Report Server contains an authorization bypass by spoofing vulnerability that allows an attacker to obtain unauthorized access.
CVE-2024-26169	Microsoft	Windows	Microsoft Windows Error Reporting Service Improper Privilege Management Vulnerability	13/06/2024	Microsoft Windows Error Reporting Service contains an improper privilege management vulnerability that allows a local attacker with user permissions to gain SYSTEM privileges.
CVE-2024-32896	Android	Pixel	Android Pixel Privilege Escalation Vulnerability	13/06/2024	Android Pixel contains an unspecified vulnerability in the firmware that allows for privilege escalation.
CVE-2020-13965	Roundcube	Webmail	Roundcube Webmail Cross-Site Scripting (XSS) Vulnerability	26/06/2024	Roundcube Webmail contains a cross-site scripting (XSS) vulnerability that allows a remote attacker to manipulate data via a malicious XML attachment.
CVE-2022-2586	Linux	Kernel	Linux Kernel Use-After-Free Vulnerability	26/06/2024	Linux Kernel contains a use-after-free vulnerability in the nft_object, allowing local attackers to escalate privileges.
CVE-2022-24816	GeoSolutionsGroup	JAI-EXT	GeoSolutionsGroup JAI-EXT Code Injection Vulnerability	26/06/2024	GeoSolutionsGroup JAI-EXT, a component of GeoSolutions GeoServer, contains a code injection vulnerability that, when programs use jt-jiffle and allow Jiffle script to be provided via network request, could allow remote code execution.

6 Research highlights

6.1 Mass exploitation: The vulnerable edge of enterprise security

A report by Stephen Robinson exploring the trend of mass exploitation of edge services and infrastructure devices, putting forward several theories as to why they have been so heavily and successfully targeted by attackers, along with examples of incidents and campaigns targeting them.

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