
# It takes two: The 2025 Sophos Active Adversary Report

**The dawn of our fifth year deepens our understanding of the enemies at the gate, and some tensions inside it; plus, an anniversary gift from us to you**

By John Shier, Angela Gunn, and Hilary Wood

The Sophos Active Adversary Report celebrates its fifth anniversary this year. The report grew out of a simple question: What happens *after* attackers breach a company? Knowing the adversary's playbook, after all, helps defenders better battle an active attack. (There's a reason we started life as "The Active Adversary Playbook.") At the same time we were discussing ways to instrument a testing environment to answer that what-happens question, Sophos was preparing to launch an incident response (IR) service. A cross-team project was born.

For five years, we've presented our data – first solely from the IR service, but eventually expanding to include data from IR's sister team supporting current MDR customers -- and provided analysis on what we think it means. As we continue to refine our process for collecting and analyzing the data, this report will focus on some key observations and analysis – and, to celebrate a half-decade of this work, we're giving the world access to our 2024 dataset, in hope of starting broader conversations. More information on that can be found at the end of the report.

## Key takeaways

- Differences between MDR and IR findings show, quantitatively, the statistical value of active monitoring
- Compromised credentials continue to lead to initial access; MFA is essential
- Dwell time drops (again!)
- Attacker abuse of living-off-the-land binaries (LOLBins) explodes
- Remote ransomware poses a unique challenge / opportunity for actively managed systems
- Attack impacts contain lessons about potential detections

### Where the data comes from

As with our previous Active Adversary Report, data for this edition is drawn from selected cases handled in 2024 by two Sophos teams: a) the Sophos Incident Response (IR) team, and b) the response team that handles critical cases occurring among our Managed Detection and Response (MDR) customers. (For convenience, we refer to the two in this report as IR and MDR.) Where appropriate, we compare findings from the 413 cases selected for this report with data from previous Sophos X-Ops casework, stretching back to the launch of our IR service in 2020.

For this report, 84% of the dataset was derived from organizations with fewer than 1000 employees. This is lower than the 88% in our previous report; the difference is primarily (but

not entirely) due to the addition of MDR's cases to the mix. Just over half (53%) of organizations requiring our assistance have 250 employees or fewer.

And what do these organizations do? As has been the case in our Active Adversary Reports since we began, the manufacturing sector was the most likely to request Sophos X-Ops response services, though the percentage of customers hailing from Manufacturing decreased from 25% in 2023 to 16% in 2024. Education (10%), Construction (8%), Information Technology (7%), and Healthcare (6%) round out the top five. In total, 32 industry sectors are represented in this dataset.

Further notes on the data and methodology used to select cases for this report can be found in the Appendix. SecureWorks incident response data is not included in this report.

### IR and MDR: What's the difference?

Though both of the datasets we use are derived from response activity, there is a critical difference in how they are generated. IR data comes from customers who come to us without MDR services already in place; they may call us when they suspect an incident is underway, or they may simply be referred by their insurance company or otherwise familiar with Sophos. MDR data comes from current managed customers (so, customers with at least some Sophos monitoring and logging services in place) who need incident response to neutralize active threats and remediate the actions of attackers; in almost all cases, we initiate notice to them that something bad is happening.

## The main event: MDR vs IR

As we compiled and normalized the IR and MDR datasets, the Active Adversary team hypothesized that we would likely observe better security outcomes in organizations where skilled active monitoring and logging were already in place – in other words, the MDR cases. While that may seem obvious, it's the magnitude of some of the differences that surprised us, and it is those differences we'll highlight in this report.

### We're one (but we're not the same): Ransomware and dwell time

In the previous report cycle, we observed, but did not report on, distinct differences between the attack types prevalent for MDR customers and

those prevalent for IR customers. This was the first strong indication of the gap between the two datasets, and it was that difference which set the tone and focus for this report.

In all previous reports, ransomware has dominated the charts, as one might expect from IR-derived data. A ransomware attack is simply too damaging for many organizations to remediate on their own, especially smaller organizations that may lack the resources necessary to mount a full response.

The previous four years of IR-only data saw ransomware occurrence vary between 68% and 81% of cases. For 2024 it is down to 40% of cases, losing its top spot to network breaches at 47%. When we break it down by data origin, the proportion for IR cases looks very much like all previous data. Ransomware (65%) is the dominant attack type, followed by network breaches (27%). The MDR data paints a different picture, in which network breaches (56%) outpace ransomware (29%) almost two to one.

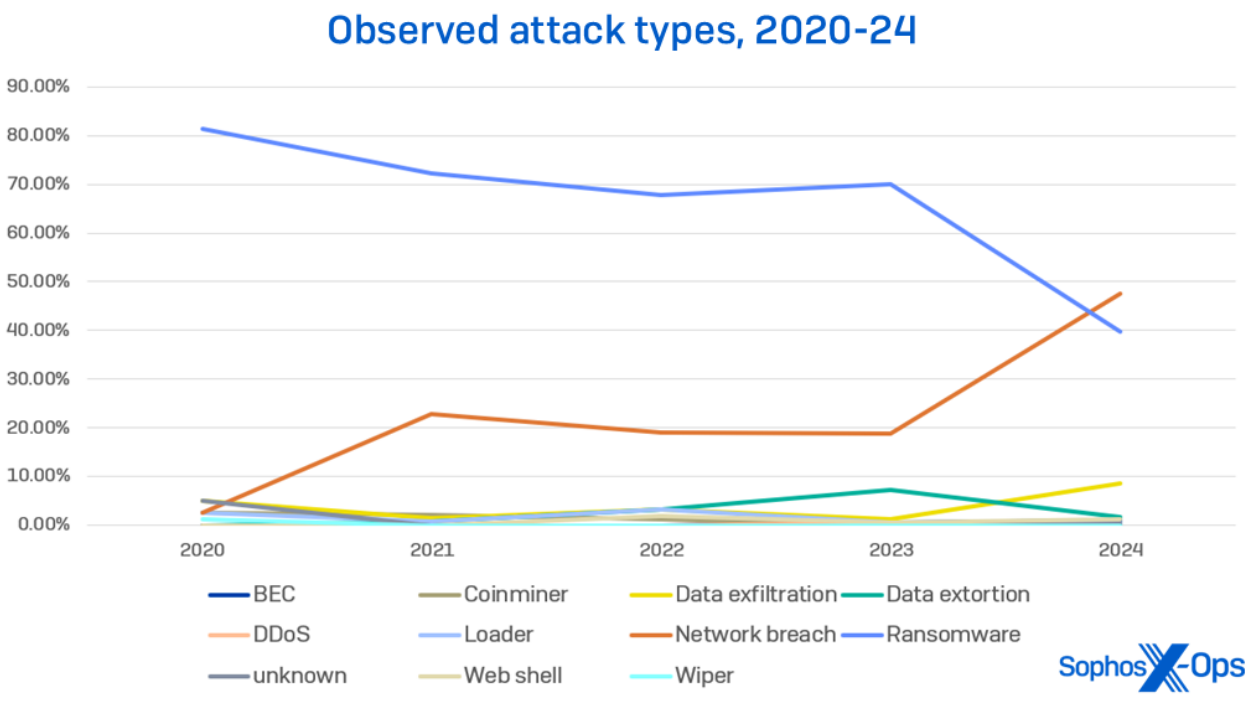


*Figure 1: The change in attack-type findings in our dataset is striking – in 2024, network breaches overtook ransomware as the attack type we most commonly observed. At the bottom of the chart, however, there's another remarkable story – whatever the dataset, whenever the year,* no *attack type rises above 10 percent of all cases*

*seen; whether ransomware or network breaches are the main event in a given year, everything else is frankly secondary*

The second set of data supporting our hypothesis concerns dwell time. Previous years have seen dwell time decreasing but stabilizing in the last few reports. (We treated dwell time to a deep analysis in our 1H 2024 report.) As far as we were concerned, dwell time was dead -- until we saw the statistics for this year.

We won't bury the lede: Median dwell time for all cases in 2024 was a swift two days. We see a familiar pattern emerge in IR cases: Overall median dwell time is 7 days, with ransomware cases at 4 days and non-ransomware cases at 11.5 days. MDR dwell times, on the other hand, were lower across the board, and the order of dwell times for ransomware (3 days) and non-ransomware (1 day) attacks were inverted.

We believe this is because certain actions (for instance, exfiltrating the data) cannot go any faster, since they rely on human activity, data throughput, or other fairly rigid time frames. That's not to say the attacks can't be done faster, because they can, but the data shows that ransomware attacks have traditionally required longer timeframes than other attack types. The fact that dwell times for ransomware cases handled by each service were roughly equal is therefore not surprising.

Non-ransomware cases, on the other hand, have fewer speed bumps, and here's where the data highlights the differences between the services. For example, with IR cases, an attackers may reside in the victim's network undetected for much longer, until an event occurs that causes sufficient noise or impact. An attacker using valid credentials, who silently exfiltrates data from a network over expected channels, might not be detected until they contact the victim, if they ever do. (It should also be noted that the ransomware sector has attracted a great many of the more amateurish type of attacker, which is usually less adept at keeping quiet and covering its tracks. Ransomware is still a numbers game, so getting knocked off a high percentage of systems is just part of the business model.)

MDR cases for non-ransomware (or pre-ransomware) incidents, on the other hand, are generated more quickly due to a combination of detection engineering and constant vigilance. Suspicious events are investigated sooner, and those that warrant additional investigation are escalated. In short, faster detection often leads to aborted ransomware, which means a higher proportion of attacks classified as network breaches -- and better outcomes for the victims.

## Come together: Root cause

In contrast, we didn't see much difference between IR and MDR cases when it came to root causes. Here we see the familiar combination of compromised credentials (41%) and exploiting vulnerabilities (22%) leading the way once again, and brute force attacks (21%) muscling their way to third place, as shown in Figure 2.

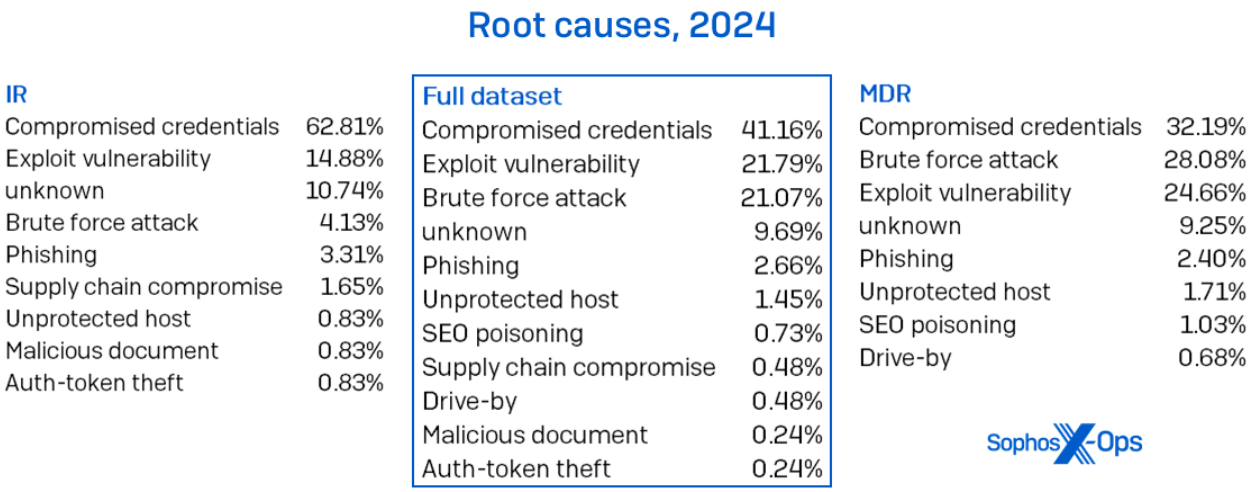


*Figure 2: Root cause in 2024 varied between MDR and IR cases, but compromised credentials are still the leading cause of pain in both datasets*

Brute force attacks have been perennially relegated to the also-ran category in the IR data, but saw a dramatic increase in the MDR data, which vaulted the attack type up the rankings for 2024. This may be down to a difference in the available root-cause data. In IR investigations, logs are often unavailable, which reduces the investigative team's ability to determine the root causes of the attack. In contrast, MDR investigations have more consistent data sources available, which allows for more precise analyses.

A look at the year-to-year data, as shown in Figure 3, shows the change in percentages between previous years and 2024.

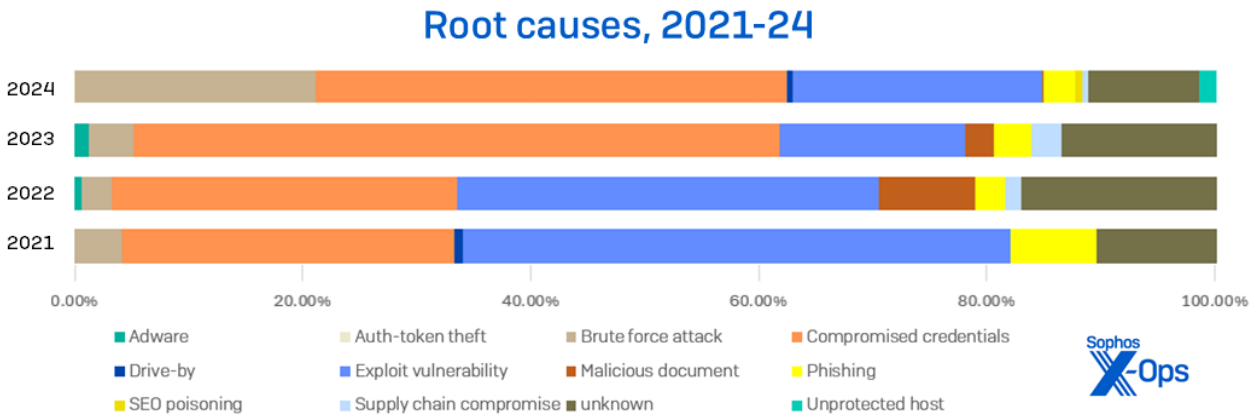


*Figure 3: Compromised credentials in 2024 retreated from previous high levels as the most common root cause of problems, but it's still a bad situation. (Data from 2020 cases is not represented in this chart due to a change in our data labeling for this category)*

In 2024, logs were missing in 47% of cases – 66% for IR, 39% for MDR. The leading reason for missing logs in all cases was that they were simply unavailable (20%) to analysts during the investigation, followed by 17% of logs being cleared by the attackers and 7% missing due to insufficient retention periods.

(One tool that often gets used to clear logs is the Microsoft binary wevtutil.exe [the Windows Event Utility]. This will generate Windows event log IDs 1102 [for security logs] and 104 [for system logs]. Organizations should consider configuring their security tools and threat hunts to detect this activity.)

The rise in brute force as a root cause aligns well with initial access (TA0001) statistics. External Remote Services (T1133) was the favored initial access method, observed in 71% of cases. As we've stated previously, this is often tightly coupled with Valid Accounts (T1078); this year the duo teamed up in 78% of cases. Exploiting a Public-Facing Application (T1190) was the second-most single contributor to initial access. The top vulnerability directly exploited for initial access was CVE-2023-4966 (Citrix Bleed; 5%). Other factors included

exposed Remote Desktop infrastructure (18%), vulnerable VPNs (12%), and exposed internal services (11%).

## You down with TTP?

We demonstrated in a previous report that there were few differences in TTPs between attacks with short (5 days or fewer) versus long (more than 5 days) dwell times. Those data were exclusively IR cases. Looking at the TTPs from this year's report, we see the pattern hold when comparing IR and MDR cases.

There were slightly more artifacts seen in MDR cases (+24%), though the MDR dataset was around 240% larger than that taken from IR. There was a 60% overlap in the 10 tools most used by attackers. Among the top legitimate tools being abused were some familiar names: SoftPerfect Network Scanner, AnyDesk, WinRAR, and Advanced IP Scanner, as shown in Figure 4.

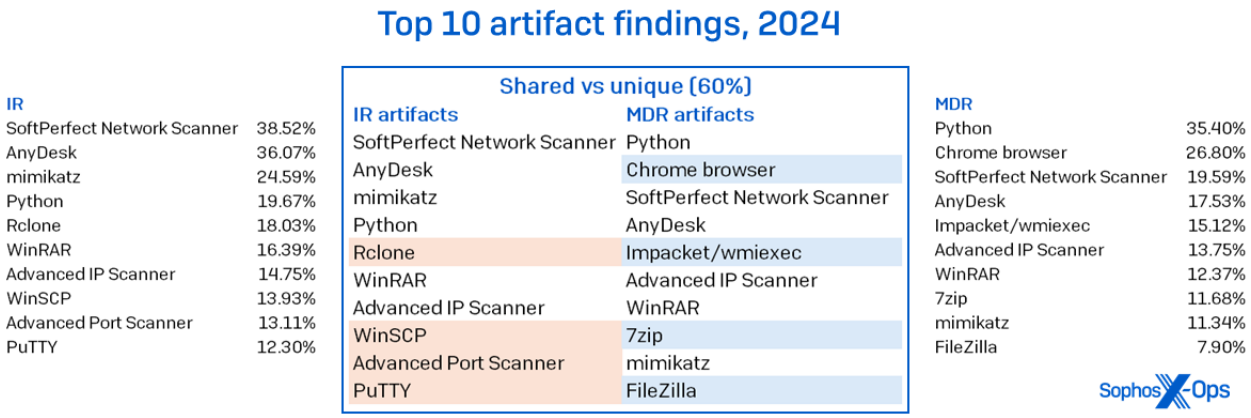


*Figure 4: The tools seen abused in IR and MDR cases didn't vary much at the top of the charts, but certain differences and absences are striking*

Microsoft binaries exhibited a tighter correlation between the datasets. The top 10 abused LOLBins had a 70% overlap, as shown in Figure 5. There was a slight shuffle in the top spot, with cmd.exe beating out RDP as the most abused LOLBin in the MDR case load. This isn't entirely surprising, since many MDR cases have a limited blast radius: When authorized to do

so, analysts will automatically isolate affected hosts, thereby limiting attackers' lateral-movement capabilities.

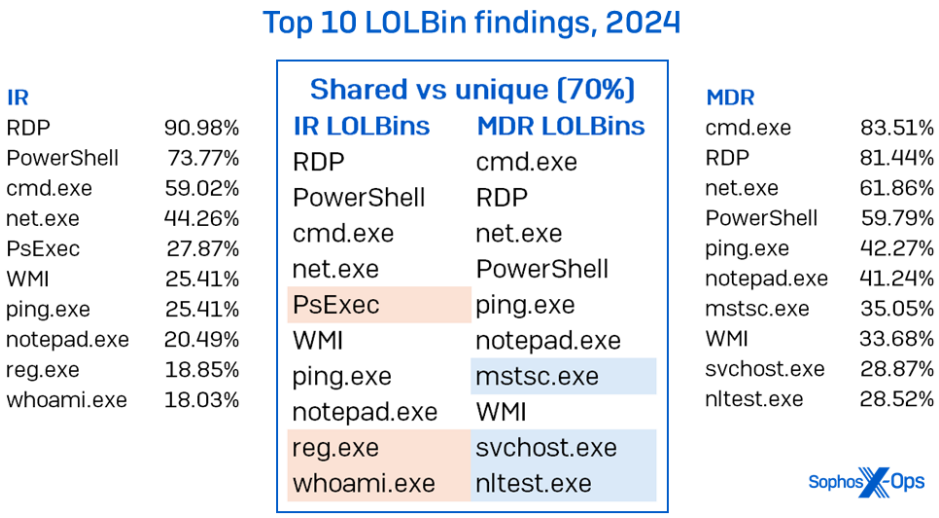


*Figure 5: LOLBin abuse presents itself much the same no matter which team is looking; in particular, the difference between MDR and IR when it comes to RDP abuse exists but is not substantial*

The final comparison looks at the "other" category, in which we group techniques and traces that don't fall into the other two categories. The top 10 had an 80% overlap in IR and MDR cases; creating accounts, deleting files, installing services, malicious scripts, and modifying the registry were the dominant techniques, as shown in Figure 6. Others, such as SAM (Security Account Manager) dumping, were more common in one team's dataset.

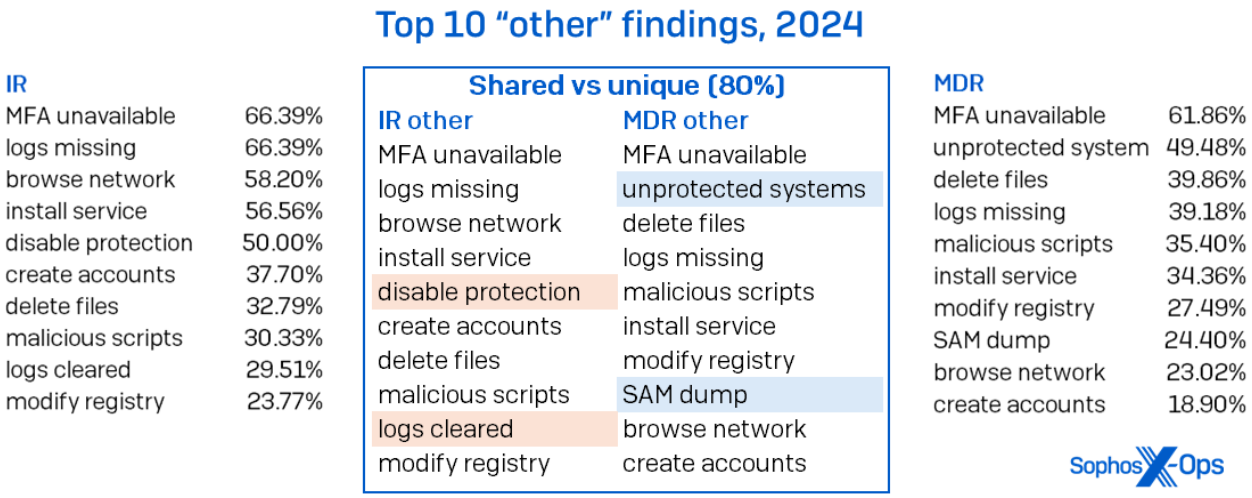


*Figure 6: As we see, in more than half of all cases, the attackers used familiar and similar TTPs. (Note that percentages add up to over 100%, since most cases have multiple findings in this category)*

## The bite from inside (reprise)

As has become the norm at Active Adversary Report HQ, we like to check in on some of our findings from previous reports, especially those for which the data period was less than 12 months. The next section looks at the key takeaways from our previous report (covering the first six months of 2024) and compares them to the full year's dataset.

### LOLBins

The abuse of Microsoft binaries continued unabated in the second half of 2024, and the ratio of unique LOLBins to previous years also continued to rise. In the first half of 2024 we saw a 51% rise in the count of unique LOLBins, which finished the year at 126% over 2023 counts. There was a 17% case rise in 2H 2024 and a 24% rise in unique binaries used. There were no meaningful differences in the individual binaries used throughout the year. Between the first half and second half of the year, there was a 95% overlap in the 20 most-abused tools in IR and MDR cases. Tools that can be used for enumeration – in addition to legitimate and malicious uses -- continued to be highly represented in both datasets, making up 50% of the 20 most-abused binaries.

Notepad.exe was a new entry in this year's top 10. This tool was predominantly used for browsing files on the network, including files containing passwords stored in plaintext (5%). Tools like Notepad provide an interesting detection opportunity. We would argue that most users are not using Notepad in favor of other Office programs. But there's also a big difference between clicking on the Notepad icon, typing notepad in Windows search, or typing notepad.exe at the command line. Being able to discriminate between these three different launch methods can inform the intent of its use.

The same is true of tools like PowerShell. We're not going to suggest that IT teams stop using it, but there are some quick heuristics that can be applied using detection engineering. Was that PowerShell script heavily obfuscated, and did it reach out to the internet? If it did, it should probably be investigated.

The main issue with LOLBins is they tend to generate a lot of noise. The challenge for IT teams is understanding where the signal exists.

## RDP

RDP detections continue to top the chart of abused Microsoft tools. In 2024, it was used by attackers in 84% of cases, with 67% being used only for internal lateral movement and 3% being used only externally. That's before we add the cases where it was used both internally and externally. The addition of those cases brings the totals to 83% and 19% respectively.

Despite RDP's continued abuse – and our pleas for it to be banished beyond the wall – we understand why it persists in networks. To that end, it provides us with an opportunity to explore how we might both constrain its use and instrument some detections for its abuse.

Ideally, all RDP use is constrained by both network choke points and user identities. Where possible we need to add MFA to the authentication flow and apply the principle of least privilege. By constraining its use, and understanding what normal looks like, it becomes easier to detect anomalies.

There are multiple ways to detect authentication events, but broadly speaking, you can look for Windows logging event IDs 4624 and 4625. The former is a successful authentication event, while the latter indicates a failed attempt. Successful login events can help you catch an attacker using valid credentials outside of normal use, while multiple failed attempts can give you an early warning to any brute force activity against your accounts.

If you use a corporate standard for naming your devices, as many companies do, you can use that as another indicator. Any successful authentication that does not conform to the standard should be investigated. If your organization does not have a standard, this could be an opportunity to implement one and create passive trip wires for attackers. Then again, if the hostname "kali" shows up on your network, as it did in 6% of cases, you should investigate.

Finally, you can take advantage of time-zone bias in RDP logging. This is the remote client's time offset from UTC. If most of your users are in UTC-6, but an otherwise-unremarkable remote client logs in using valid credentials and a normal looking hostname, but has a time-zone bias of +3, run like hell to find out why. (And then there are the times we've seen innocuous-looking machines connected, but sharing a Russian-named printer for some reason…)

The idea behind these detection opportunities is to take independent, but sometimes noisy or weak signals, and stitch them together to achieve a stronger, more reliable signal. Or, as the cool kids call it, *defense in depth*. Those wanting to know more about RDP and how to detect its abuse can find additional details in our RDP series.

## Attribution

In the last report, we predicted that in 2024 there would ultimately be no overwhelmingly dominant ransomware adversary; with a law enforcement takedown early in the year kneecapping LockBit, 2023's leading miscreant, the field opened up for the Next Big (Bad) Thing. As the table in Figure 7 shows, this was correct – Akira rose to the top of the pack, but only just. (LockBit was, on the other hand, so dominant at the beginning of last year that it still came in third in the rankings despite the takedown.) During the second half of the year, Fog seeped onto the charts, edging out Akira for the top spot. (The MDR team did see a couple of trailing-edge LockBit infections early in the second half, but even those traces evaporated by year's end.) The pattern may yet break down in 2025 thanks to likely changes in (among other things) law-enforcement effort coordination – and LockBit still swears they're making a comeback. We'll be watching with interest.

## 2024 ransomware attributions (prevalence >2%)

| 2024 H1 | | 2024 H2 | | 2024 FY | |
|---|---|---|---|---|---|
| LockBit | 21.54% | Fog | 21.21% | Akira | 14.63% |
| Akira | 9.23% | Akira | 18.18% | Fog | 12.80% |
| CryTOX | 6.15% | RansomHub | 9.09% | LockBit | 11.59% |
| Faust | 6.15% | Play | 5.05% | RansomHub | 6.71% |
| Qilin | 6.15% | LockBit | 5.05% | Black Suit | 4.88% |
| Black Suit | 6.15% | Black Basta | 4.04% | Black Basta | 4.27% |
| Black Basta | 4.62% | Black Suit | 4.04% | Play | 3.66% |
| Hunters Intl | 4.62% | Abyss | 2.02% | Qilin | 3.66% |
| RansomHub | 3.08% | Lynx | 2.02% | CryTOX | 3.66% |
| 8Base | 3.08% | Hunters Intl | 2.02% | Faust | 3.05% |
| ALPHV/BlackCat | 3.08% | Qilin | 2.02% | Hunters Intl | 3.05% |
| Mario | 3.08% | Cicada3301 | 2.02% | | |
| | | CryTOX | 2.02% | | |
| | | GlobeImposter | 2.02% | | |

Sophos X-Ops

*Figure 7: Fame is fleeting, as LockBit's perpetrators learned in the latter half of 2024; meanwhile, a heavy Fog rolled in*

Being able to attribute trouble to a specific adversary is soothing, somehow. But practitioners are often fighting forces that are nominally on their side, while dealing with choices made by the larger business that feel like one more conflict to be handled. Our case study for this report describes how that went for one "unlucky" MDR customer.

## Case study: Two against one

While we continue to reiterate fundamental security tenets (close exposed RDP posts, use MFA, and patch vulnerable systems), in the face of business change processes beyond practitioners' control, it's not always that easy. Security practitioners are not only fighting the battle against the threats posed by external adversaries, but an internal struggle with business processes and change management. This tug-of-war came back to bite one MDR customer. Following a network breach in which the threat actor gained initial access through a vulnerable VPN, the customer faced a two-month estimated timeframe to patch the VPN appliance. With a ransomware gang waiting in the wings, the conflict between security priorities and those of the larger business resolved in just about the worst way possible.

## You and me against me

The Sophos MDR team recently responded to this customer's critical incident, with initial access identified as one of our usual suspects – an unpatched VPN appliance. In this case, a FortiGate firewall was running on firmware version 5.6.11, which was released in July 2010; the firewall itself reached end-of-life in October 2021. In addition, MDR identified a misconfiguration in VPN user-access controls, which significantly increased the risk of unauthorized access.

After gaining initial access, the threat actor moved laterally to the domain controller, leveraged AV-killer tools, performed enumeration, and gained persistence on a number of devices within the estate. At this stage, MDR's response team disrupted the attacker activity, and calm resumed.

The MDR team recommended the customer (at minimum) patch the 14-year-old VPN firmware with urgency, and disable the SSL VPN in the meantime. However, the customer's business processes were not cooperative; disabling the VPN altogether would cause unacceptable business impact, and the patches couldn't be applied for two months (!). The misconfiguration, the customer estimated, would take one week to remedy.

## Already fighting

It's an unfortunate fact of incident-response life that we cannot compel; we can only recommend – and, sometimes, we can only stand by watching history repeat itself. And it *was* repeating: The same customer had already experienced a similar breach, involving the same vulnerable VPN, 14 months earlier. In that case, the customer did not yet have MFA enabled for VPN logins; a brute force attack was successful, and the attacker was able to disable protections and dump credentials. In the process, the attacker managed to compromise a key service account, leaving the customer unable to perform a crucial credential reset due to – again – business requirements. (Remember that service account; we're about to see it again.)

The gap between the first breach and the second was, as mentioned, 14 months. The gap between the second and the third was far shorter.

### So what's another one?

The second incident concluded. The VPN and that service account – one thing out of support for nearly four years, one thing known-compromised for over a year – waited in business-process limbo, as did the VPN misconfiguration. The security practitioners were patient. The attacker wasn't. Nine days after the close of the second breach, CryTOX roared in. Using the compromised service account and taking full advantage of the unpatched and (still) misconfigured VPN, the ransomware ran rampant through the system, moving laterally, killing endpoint-security processes, and ultimately encrypting the entire estate.

It may be said in this case that ransomware won the tug of war between security practices and business change processes. (Silver lining: After the third incident, the VPN was finally disabled, though affected accounts were still re-enabled without credential resets.) While not all organizations are so unlucky, in this case the wait for business change approval was a risk-assessment gamble that failed terribly.

## Best of the rest

As we wrap up our 2024 findings, let's check in on other statistics that drew our attention.

In addition to an increased number of cases, this year's dataset included the biggest year-to-year increase in all observed TTPs. In comparison with 2023, the number of abused tools was up 80%, LOLBins were up 126%, and everything else ("other") was up 28%. What's interesting about these numbers is the long tail for each category – that is, the number of tools or LOLBins or "other" that appeared ten times or fewer in the dataset. When we tally every single finding in every single case, those rarities account for 35% of all tool use (689 findings of 1945 total; 334 unique items), 12% of all LOLBin use (508 findings of 4357; 184 unique items), and 12% of all "other" (476 findings of 4036; 189 unique items). A biologist might call those

vestigial tails; we call them a lower investigation priority than the dominant beasts at the tops of the TTP charts.

## No time to waste

When it comes to certain objectives, attackers don't fritter and waste the hours in an offhand way. We first reported on the race to Active Directory compromise in 2023. This statistic has continued to trend downward, and the median now stands at 0.46 days. In other words, once an attacker enters the environment, it's only 11 hours before they go after the AD server. Most (62%) of the compromised servers were running operating systems that were out of mainstream support.

## Games without frontiers

Another time-related statistic that we first reported on in 2023 was the time of day that attackers chose to deploy ransomware payloads. While more data softens the values somewhat, the results are still compelling. In 2024, 83% of ransomware binaries were deployed outside the target's local business hours; the all-time statistic stands at 88%. While it appears that ransomware deployments only come out at night, there does not however seem to be any lingering preference in days of the week.

## Tools to walk through life

The proportion and types of tools – both legitimate and malicious – that make up this category have remained relatively stable for many years. Here are some highlights from this year's data, in addition to the issues covered above.

We've seen a big drop in the proportion of attacks that use Cobalt Strike. This tool occupied the top spot in abused tools from 2020-2022, dropping to second place in 2023. This year saw it slip all the way down to thirteenth on our list, appearing in just 7.51% of cases. Due to its historical popularity with attackers, it still occupies the top spot in the all-time rankings, where it has been involved in 25% of attacks in the past five years. We believe the decrease is due to increased prevention and detection capabilities. Cobalt Strike was popular because it

was effective. Now that its effectiveness has declined, so has its use. While this is welcome news, it also suggests that something else has or will take its place.

A tool that has seen an order of magnitude increase in abuse is Impacket. Impacket tools have been around for at least a decade and can perform a variety of actions, including manipulating network protocols, dumping credentials, and reconnaissance. Its use has steadily grown in recent years, from 0.69% in 2021 to 21.43% in 2023; attackers really ramped up their use of Impacket in 2024, when it overtook all other tools and landed in the top spot. The most used Impacket tool was wmiexec.py, which featured in 35% of attacks. (In our statistics, we identify the specific Impacket subclass whenever possible; if there is doubt, we simply classify it as Impacket, no subclass.)

A venerable tool seeing a slight year-on-year decline is mimikatz. The credential-harvesting tool was reliably observed in around a quarter of attacks in previous years but slipped to 15% in 2024. While we can't decisively attribute its decline to any one thing, it's possible that it is related to the increased use of Impacket tools; specifically, the secretsdump.py script that can be used to dump hashes from remote machines. This correlates with a year-on-year increase in remote registry dumping and a halving of LSASS dumps (most commonly attributed to mimikatz in our data). Secretsdump.py was seen in at least 6% of attacks and was the second most used Impacket tool after wmiexec.py.

Of the top 15 tools being abused, 47% are often used for exfiltration of data. These tools include well-known archiving software and file transfer tools.

## Other findings

Since we started tracking the availability of multifactor authentication (MFA) in breached organizations, the news has gotten worse. In 2022, we observed 22% of victims did not have MFA configured. That proportion nearly tripled to 63% in 2024. This is one area where there was no meaningful distinction between IR and MDR cases. MFA was unavailable in 66% of IR

cases and 62% of MDR cases. This highlights one way in which even the most capable detection and response program can still leave organizations vulnerable to attack.

Another concerning metric was the proportion of unprotected systems found in breached organizations. In 40% of the cases we investigated, there were unprotected systems. When we consider there were also vulnerable VPNs (12%), vulnerable systems (11%), and end-of-life systems (5%) in some of these environments (this report's case study, for instance, had all three), attackers might feel like a cunning fox in the chicken's lair.

Some may ask why we're still seeing ransomware cases at all in an MDR service. One big reason has to do with unprotected systems and their relationship with remote ransomware. All that malicious activity – ingress, payload execution, and encryption – occurs on unmanaged machines, therefore bypassing the organization's security tools. The only indication of compromise is the transmission of documents to and from other machines. Our telemetry indicates that there has been a 141% year-on-year increase in intentional remote encryption attacks since 2022, as shown in Figure 8. (We've talked previously about remote ransomware and how to parry it, including a deep dive into our CryptoGuard technology; as the numbers rise, remote ransomware may be a major topic in a later Active Adversary Report.)

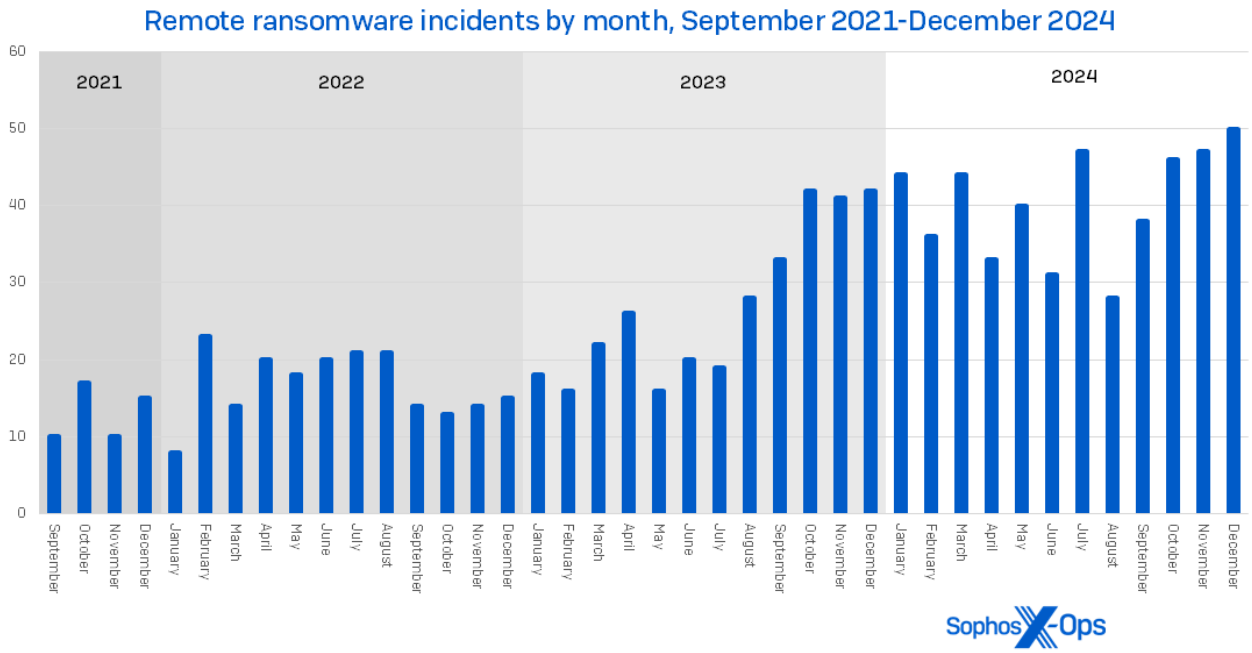


*Figure 8: According to Sophos X-Ops data, 2024's remote ransomware tally was 141% of that of 2022; note the startling rise in cases over the last 18 months of the data*

The lack of visibility for files moving around the network – and of missing logs - also contributes to exfiltration statistics. In 2024, analysts were able to confirm that exfiltration occurred in 27% of cases. When we include evidence of data staging and possible exfiltration, this rises to 36%. Ransomware victims had their data exfiltrated in 43% of the incidents we investigated. An additional 14% had possible exfiltration or evidence of data staging. Unlike time-to-AD, exfiltration findings occur towards the end of an attack. There was a median time of 72.98 hours (3.04 days) between the start of an attack and exfiltration, but only 2.7 hours (0.11 days) from exfiltration to attack detected for ransomware, data exfiltration, and data extortion cases.

## Bring the noise

Finally, this report has traditionally looked at MITRE impacts (TA0040). Given ransomware's prevalence in the data, it's not surprising that Data Encrypted for Impact (T1486) tops the chart, as it has every year. But looking at the rest of the impacts, we see an opportunity for defenders: The causes of many of the other impacts are events that can be detected.

## Attack impact (percentage of cases), 2020-24

| 2020-23 (IR only) | | 2024 (IR + MDR) | | All-Time (2020-24) | |
|---|---|---|---|---|---|
| Data Encrypted for Impact | 71.75% | Data Encrypted for Impact | 39.23% | Data Encrypted for Impact | 57.52% |
| No Impact | 18.64% | Data Manipulation | 37.77% | Data Manipulation | 25.85% |
| Data Manipulation | 16.57% | No Impact | 30.99% | No impact | 24.05% |
| Inhibit System Recovery | 13.37% | Account Access Removal | 17.43% | Inhibit System Recovery | 13.56% |
| Account Access Removal | 9.04% | Inhibit System Recovery | 13.80% | Account Access Removal | 12.71% |
| Resource Hijacking | 7.16% | Service Stop | 11.86% | Service Stop | 6.25% |
| System Shutdown/Reboot | 5.08% | System Shutdown/Reboot | 3.87% | System Shutdown/Reboot | 4.56% |
| Financial Theft | 3.77% | Financial Theft | 1.21% | Resource Hijacking | 4.45% |
| Service Stop | 1.88% | Resource Hijacking | 0.97% | Financial Theft | 2.65% |
| Data Destruction | 0.38% | Data Destruction | 0.73% | Data Destruction | 0.53% |
| Network Denial of Service | 0.19% | | | Disk Wipe | 0.11% |
| Disk Wipe | 0.19% | | | Network Denial of Service | 0.11% |

Sophos X-Ops

*Figure 9: MITRE's Impact categories change over time, but Data Encrypted for Impact's reign at the top of the Active Adversary charts is unbroken throughout our five-year history, including both IR's and MDR's cases this year. (Note that percentages add up to over 100%, since some cases have multiple impacts)*

For instance, Inhibit System Recovery (T1490) is often invoked because the threat actor deleted volume shadow copies. Tools like vssadmin.exe, the shadow-copy management tool (seen abused in 10% of all cases), or the WMI command line (seen abused in 24%) are used to do the deed. You can also detect when vssadmin is used to create shadow copies, which precedes its exfiltration. Likewise, we saw attackers delete files in 26% of all cases. In that circumstance, watching for unexpected use of del.exe may be a sign of adversary action. Detection engineering can listen for suspicious events of this ilk, to hear the noise attackers make when they're trying to cause you harm.

## Conclusion

To the practitioners out there, we see you. You're doing the work and you know the business. You also know the limitations of what you can accomplish. The good news is that you don't need to be helplessly hoping things will get better, especially when help is available.

To the business and tech leaders, give your teams a chance. We know money and resources are tight. That often means loading up your IT staff with more work and responsibility than they can handle. Though it may sound self-serving coming from a research team attached to

a security vendor, we believe IT teams need to focus on how they enable the business and let experts do the dirty work of fighting the attackers. Because one thing is clear from the data: When there's someone paying attention to the environment and they are able to act quickly and decisively, outcomes dramatically improve. The alternative is repeating mistakes from the past. The choice is yours: You can get with this, or you can get with that. We think you'll get with this, for this is where it's at.

### Acknowledgements

The authors wish to thank the Sophos IR and MDR teams, Mark Loman, Chester Wisniewski, and Matt Wixey for their contributions to the AAR process.

## Appendix: Demographics and methodology

For this report, we focused on 413 cases that could be meaningfully parsed for information on the state of the adversary landscape throughout 2024. Protecting the confidential relationship between Sophos and our customers is of course our first priority, and the data herein has been vetted at multiple stages during this process to ensure that no single customer is identifiable through this data – and that no single customer's data skews the aggregate inappropriately. When in doubt about a specific case, we excluded that customer's data from the dataset.

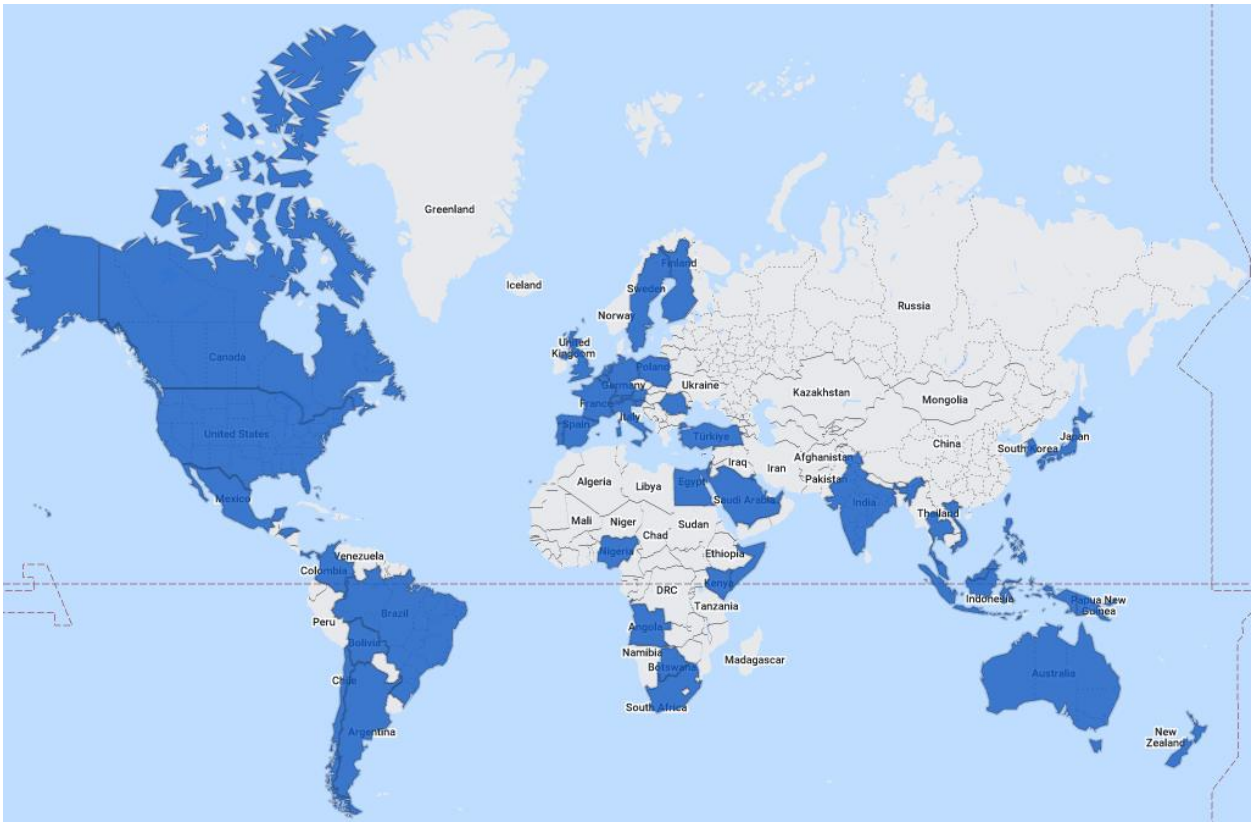


*Figure A1: We get around: It's Sophos Incident Response and MDR at work around the world (map generated courtesy of 29travels.com)*

The following 57 nations and other locations are represented in the full dataset:

| | | |
|---|---|---|
| Angola | Hong Kong | Qatar |
| Argentina | India | Romania |
| Aruba | Indonesia | Saudi Arabia |
| Australia | Israel | Singapore |
| Austria | Italy | Slovenia |
| Bahamas | Jamaica | Somalia |
| Bahrain | Japan | South Africa |
| Belgium | Kenya | South Korea |
| Bolivia | Kuwait | Spain |
| Botswana | Malaysia | Sweden |
| Brazil | Mexico | Switzerland |

| | | |
|---|---|---|
| Canada | Netherlands | Taiwan |
| Chile | New Zealand | Thailand |
| Colombia | Nigeria | Turkey |
| Egypt | Panama | Turks and Caicos Islands |
| Finland | Papua New Guinea | United Arab Emirates |
| France | Philippines | United Kingdom |
| Germany | Poland | United States of America |
| Honduras | Portugal | Vietnam |

### Industries

The following 32 industries are represented in the full dataset:

| | | |
|---|---|---|
| Advertising | Financial | News Media |
| Agriculture | Food | Non-profit |
| Architecture | Government | Pharmaceutical |
| Communication | Healthcare | Real estate |
| Construction | Hospitality | Retail |
| Education | Information Technology | Services |
| Electronics | Legal | Transportation |
| Energy | Logistics | Travel and tourism |
| Engineering | Manufacturing | Utilities |
| Entertainment | Mining | Wholesale |
| Finance Services | MSP/Hosting | |

### Methodology

The data in this report was captured over the course of individual investigations undertaken by Sophos' X-Ops Incident Response and MDR teams. For this first report of 2025, we gathered case information on all investigations undertaken by the teams throughout 2024

and normalized it across 52 fields, examining each case to ensure that the data available was appropriate in detail and scope for aggregate reporting as defined by the focus of the proposed report. We further worked to normalize the data between our MDR and IR reporting processes.

When data was unclear or unavailable, the authors worked with individual IR and MDR case leads to clear up questions or confusion. Incidents that could not be clarified sufficiently for the purpose of the report, or about which we concluded that inclusion risked exposure or other potential harm to the Sophos-client relationship, were set aside. We then dissected each remaining case's timeline to gain further clarity on such matters as initial ingress, dwell time, exfiltration, and so forth. We retained 413 cases, and those are the foundation of the report. The data offered in the downloadable dataset has been further redacted to ensure customer confidentiality.

About the Author

### John Shier

John Shier is a Field CTO at Sophos. John is a popular presenter at security events, and is well-known for the clarity of his advice, even on the most complex security topics. John doesn't just talk the talk: he also gives hands-on technical support and product education to Sophos partners and customers.

About the Author

### Angela Gunn

Angela Gunn is a senior threat researcher in Sophos X-Ops. As a journalist and columnist for two decades, her outlets included USA Today, PC Magazine, Computerworld, and Yahoo Internet Life. Since morphing into a full-time technologist, she has focused on incident response, privacy, threat modeling, GRC, OSINT, and security training at companies including Microsoft, HPE, BAE AI, and SilverSky.

About the Author

### Hilary Wood

Hilary Wood is a Senior Threat Analyst in the Sophos Managed Detection and Response (MDR) Team, working closely with MDR customers to respond to critical security incidents and mitigate evolving threats. Hilary is passionate about analyzing trends across the cyber threat landscape in order to assist organizations in staying ahead of both persistent and emerging cyber threats.