Nowhere, man: The 2026 Active Adversary Report

AI headline hype didn't deliver a sea change for practical defense — but one below-the-radar development should

February 24, 2026

Written by John Shier, Hilary Wood, Angela Gunn

Security Operations | Threat Research | Active Adversary | Active Adversary Report

SHARE THIS

In a world where so much changes rapidly, it can be interesting and informative to identify when things stay the same. Throughout 2025 many people claimed — as they have for a couple of years now — that *this* was going to be the year in which AI was going to make a meaningful difference in the threat landscape.

Aside from some provable uses of AI to supercharge phishing and other social scams, and a fair number of overdramatic headlines, it just didn't happen. This year's Active Adversary Report details what happened instead — including a change that *does* demand your attention.

This year's crop of attackers used the same tools, techniques, and procedures (TTPs) they have for years. Abuse of legitimate tools remained consistent, as did the lack of blocking categories of tools that are known to be routinely abused. Missing telemetry continued to make it difficult for blue teamers to spot the signal in the noise, and an ongoing lack of phishing-resistant multifactor authentication (MFA) gave the criminals a quiet way in.

The most concerning change, meanwhile, has also been years in the making: The dominance of identity-related root causes — brute-force attacks, phishing, and other compromised-credential tactics — for successful initial access. This constellation of tactics leverages weaknesses that can't be addressed by simple patch hygiene and occasionally acts as a bonus multiplier for attacks in progress, as this year's Case Study documents.

While detection and response have proven reliable in the fight against attackers, we can't forget the role that prevention can play. This year's report recaps what we saw during incident response investigations and can serve as a guide for how to prevent some of the most common attacker TTPs. (Block Python today and thank us later. Read on to find out why.)

Key takeaways

- ✔ GenAI adds speed, volume, and noise to the threat landscape… but for now, that's about it.

- ✔ Identity-related tactics such as compromised credentials, brute-force attacks, and phishing, are by far the most common reason attackers gain initial access.
- ✔ Attackers have made few changes to specific tools, tactics, or procedures — though one weird blocking trick may make a huge difference for many enterprises.
- ✔ Saving money by minimizing telemetry collection might be penny-wise, but it's definitely pound-foolish.
- ✔ Prevention still beats detection, both in outcomes and in time and effort spent defending.

Where the data comes from

Data for this edition is drawn from selected cases handled between November 1, 2024 and October 31, 2025 by Sophos Incident Response (IR) teams, including those operating within our Secureworks group, and by 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 661 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 1,000 employees. Just over half (56%) of the 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 (19.82%) was the most likely to request Sophos X-Ops response services. Financial (8.93%), Construction (8.62%), Information Technology (6.96%), and Healthcare (6.35%) round out the top five. In total, 34 industry sectors are represented in the 2025 dataset.

Further notes on the data and methodology used to select cases for this report can be found in the appendix.

Add it up: The stats

The types of attacks our incident response services investigated, and their proportions, held steady in this year's report. Overall, we continued to see network breaches topping ransomware, driven largely by the proportion of MDR (69%) to IR (31%) cases in our data. Within each service, we again saw that proactive monitoring of the environment yielded better outcomes than reactive investigation. In this year’s roundup, we will cover some of the topline statistics and give additional insight into some interesting observations. For additional insights, please see the raw data we are once again making available on GitHub.

Our data showed a fourfold increase in business email compromise (BEC) attempts this year. We suspect that this is actually due to a new offering from the Sophos IR team: Customers are able to engage the IR team on an hourly basis, which makes it feasible for companies to launch smaller investigations based on telemetry signals. This means that organizations can get peace of mind (or at least clarity) when something suspicious occurs, without incurring the costs of a full IR investigation.

In our data, most of these engagements turned out to be compromised M365 identities leading to attempted (or, less frequently, successful) financial theft or resource hijacking, as attackers tried to use organizations’ trusted names to send out phishing emails.

Data exfiltration attacks rose to the highest percentage (12.71%) recorded since we launched our Active Adversary reports in 2021. Some of these attacks were executed by ransomware groups that had managed to exfiltrate data before their encryption attempts were thwarted. Others also stole data but did not identify themselves, and we weren't able to confidently attribute the activity to a known group. What seems clear is that many attackers will want to walk away with *something*. If they can't monetize the attack directly through ransomware, they will at least steal data in the hope of monetizing it later through extortion or additional criminal activity, which can include selling the information to other would-be attackers.

Ain’t what he seems: Root cause and identity

As mentioned above, there are important changes to root causes. The proportion of attacks that rely on compromising identity has been rising for several years — we mentioned this in last year’s AAR, though back then we thought the trend might be wavering — and now thoroughly dominates the root-cause category.

In 2025, we saw 67.32% of root causes were related to compromised identity — in our dataset, brute-force attacks, credential phishing, auth-token theft, trusted relationships, and the catchall “compromised credentials” category. (See the sidebar below for more on how the catch-all category is used.) This confirms what we and many others have been saying for a few years: Attackers aren't breaking in, they're logging in. This is also in line with the modus operandi of 2025’s Scattered LAPSUS$ Hunters group, possibly the most successful set of criminals so far to weaponize identity attacks.

About compromised credentials

We often don't get the opportunity to see all the logs — or sometimes any logs — for all the systems involved in an attack. Often, the most we can say is that the credentials were compromised. The credentials could have been stolen ages ago (for example, by initial access brokers [IAB] or an infostealer) but abused much later. Where the data supports it, we attribute the compromise to the most specific root cause. Unfortunately, too many cases simply don't have enough telemetry to determine how the credentials were compromised.

Who are you? The multi-year rise of identity failures

Even as we continue to stand by our usual recitation of infosecurity fundamentals — close exposed RDP ports, use MFA, and patch

vulnerable servers — we're always looking at how to prioritize those tasks. As noted above, a multiyear trend in changes to incident root cause -- that is, not the "how'd they get in" of our initial access statistics, but the deeper question of "why'd that initial-access attempt work?" -- indicates that attackers are currently far less reliant on unpatched vulnerabilities than they used to be.

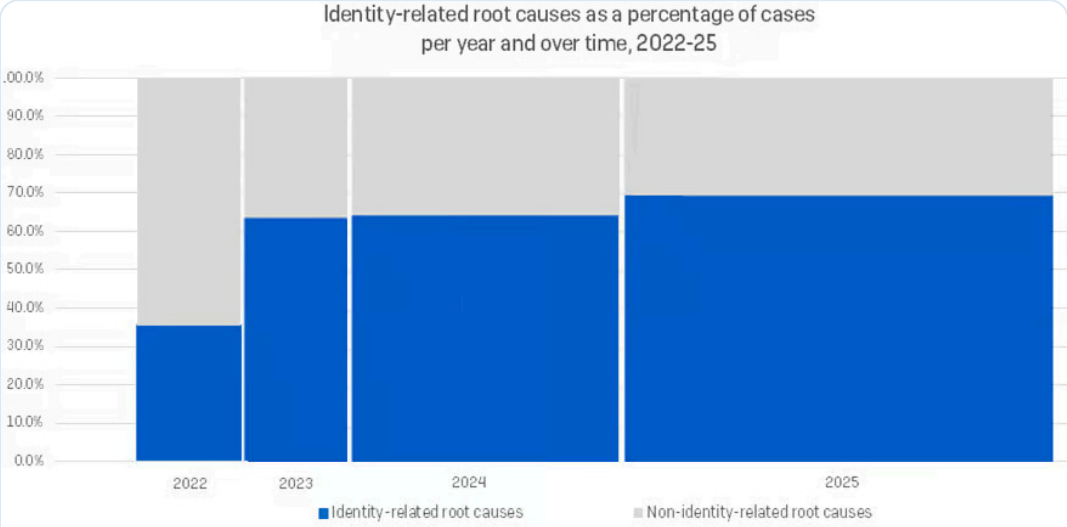

Figure 1: Identity-related root causes as a proportion of all cases covered in the last four years of the Active Adversary Report. (2022 is included by way of comparison — the last year that identity-related issues weren't the leading root cause of trouble.) The width of each bar indicates the relative number of cases we see each year. As shown above, 2025 was our heaviest case load yet and had the largest amount of identity-related root causes ever.

Unfortunately, attackers lately are far more reliant on what we collectively call "identity-related" factors. These include compromised credentials, brute-force attempts, and phishing, along with less common approaches such as trusted relationships in which privileged credentials were over-privileged, or in which access was gained by theft of an authentication token. Our 2025 data indicates that cases involving compromised credentials, with nothing further known about the means of that compromise, accounted for 42.06% of root causes. Brute force attacks (15.58% of cases), the type of identity-related compromise easiest to conclusively identify in our investigations, are nearly tied with exploitation of vulnerabilities (16.04%) as the second most common root cause.

Phishing, the third most likely identity-related root cause, accounts for just 6.35% of identifiable cases – but even that's up substantially in 2025, more than doubling its occurrence in 2024. And phishing can cause an outsized amount of mayhem when used in conjunction with

other security failures, as demonstrated in the case study later in the Report.

Just a little light: The rest of the stats

The dominance of compromised credentials and other identity-related root causes doesn't mean that unpatched vulnerabilities don't matter. Even with numbers for that root cause at their lowest since 2021, attackers seized the opportunity when it presented itself, finding the fatal error in what was otherwise all right. Based on available evidence, we were able to confirm which CVEs were used in 52 cases. (This is another interesting effect of the dominance of identity-related root causes — who needs CVEs and exploits to get in when you've got passwords?) Among the exploited vulnerability cases, over two-thirds (67.31%) exhibited use of CVE-2024-40766, the SonicOS bug repeatedly patched by SonicWall over the course of the year.

The CVEs exploited in the 2025 dataset included other zero-days as well, but also had some vulnerabilities reaching back as far as 2008. The median time between when a vendor advisory/patch was published and when it was exploited by the attacker for all confirmed exploited vulnerabilities was 322 days. Similarly, the median time between when a public proof-of-concept for a vulnerability was released and when it was exploited by an attacker was 296.50 days.

On the ransomware front, attribution remains consistent, and the most prevalent ransomware brands we saw deployed mirror those most often seen by other threat intelligence sources. Akira (in Secureworks parlance, Gold Sahara) and Qilin (Gold Feather) led the way, followed by SafePay (Gold Leapfrog), Inc (Gold Ionic), and Play (Gold Encore). These five brands made up 51% of all ransomware incidents. All are ransomware-as-a-service brands, acting as flags of convenience to the actual perpetrators of the crimes, who will align themselves with whichever brand suits their purpose at any given time, and for any given victim.

For those keeping track of the pattern in which ransomware prevalence alternates between "poplar tree" years in which no brand is prevalent,

and “banyan tree” years in which one ransomware brand overshadows all, 2025 was indeed a banyan year. The name of that banyan was Akira, with more than twice as many cases [22.58%] as the runner-up Qilin [11.06%].

In the 2025 dataset, we saw 51 unique ransomware brands deployed, with 27 brands continuing to be used from the previous year and 24 new ones debuting. Looking at ransomware brands in our dataset from 2020 to 2025, three brands (LockBit, Medusa, Phobos) and one technique (abuse of native BitLocker encryption) have persisted for the duration.

The attribution picture for data exfiltration cases is much less clear, with 84.47% of those cases having no defined attribution. Threat actors are often seen by investigators quietly entering a network, exfiltrating data, and exiting without identifying themselves or attempting contact — they take what they want and go. Some of these cases may have been classified as data extortion, but our investigators were never told if a ransom demand was made (a customer’s right, to be sure — that path is for their steps alone).

Dwell time has also stabilized at a median of three days. Curiously, this decrease is driven by both attackers and defenders. Despite attackers accelerating their activities, we also know that many defenders have made great strides in speeding up their detection capabilities in the last couple of years, and it's showing. (Keep it up!)

As with last year, there were some dwell time differences associated with case origin. All-cause IR cases (5.00 days) observed slightly longer median dwell times, despite being down 29% year-on-year, than all-cause MDR cases (2.00 days). Non-ransomware IR cases continue to generate the longest dwell times (6.00 days).

The bottom line is that, at this point, attackers are likely moving as fast as they can once they’re in the system. These days, it’s far less common for them to just poke around. Only the defenders can make this dwell-time number approach zero, and the pressure to outrun us may cause the attackers to make more noise than they might want or intend. *For now.* Automation and orchestration would be obvious tasks

for which attackers could use AI to up their game. Watch this space next year.

Let’s active: A tough year for AD

Unfortunately, the attackers' appetite for accessing Active Directory (AD) servers has not been sated. The speed with which attackers attempt to go after AD after gaining access to the system sped up by 70% over last year, down to a median of just 3.40 hours. Conversely, the time between when AD access is attempted and when the attack is detected has risen by 16%. Of the Windows Server versions we were able to identify, 13% were running end-of-life versions, with a further 27% soon to be.

Continuing our timeline observations, while looking at the entire dataset, attackers persist in deploying ransomware when organizations (in theory) aren’t looking. Translated to local times, 88.10% of ransomware was deployed during non-business hours, spread almost evenly across all days of the week, but with a slight bump on Thursdays and Fridays. Data exfiltration followed a similar trend, with 78.85% choosing to exfiltrate in off-hours. These cases showed a slight bump on Wednesdays and Thursdays. Maybe the jokers in the ransomware department were just waiting for the clowns in exfil to do their jobs?

Exfiltration for all attack types occurred a bit over three days (78.83 hours) from the start of the attack, but only 1.87 hours before the attack was detected. We observed that 49.77% of ransomware cases included confirmed exfiltration of data — over half (53.92%) when we included potential exfiltration. In most cases, a lack of firewall logs contributed to the uncertainty. Nearly half (49.07%) of the ransomware cases in which exfiltration was confirmed also had the data publicly leaked within 19.5 days of exfiltration.

On the attacker-tools, LoLBins, and otherwise-uncategorizable artifacts (“other”) front, we saw increased volume, but nothing like last year's increase over 2023 numbers. Unique tool count was up 8%, Microsoft binaries were up 19%, and everything else was up 23% year-on-year.

Some rise, some fall: Python and Cobalt Strike

Tool use year-on-year was consistent overall, but with important specific changes. This year, Impacket continues to lead the way when we combine all Impacket tools (e.g. atexec, secretsdump, wmiexec, etc.), followed closely by Python — which Impacket requires to execute. Impacket accounted for 36.01% of all tools and saw an 83.08% increase in usage over 2024.

The rise of Impacket use by cybercriminals, and its dependence on Python, provide an opportunity for defenders. Unlike PowerShell, unless Python is required by the organization, it should be summarily blocked, at least on non-development workstations, to prevent Impacket use. Careful consideration and additional controls, including appropriate logging and even ticket creation to check on suspicious Python usage, should be applied for use cases in which Python might be necessary.

In contrast to Impacket's ascendency, Cobalt Strike continues its downward trend into obscurity, barely cracking the top 35 tools in 2025. When comparing year-on-year tools use, the top 30 tools only saw five differences. Of the legitimate commercial tools being abused, AnyDesk is still the most abused remote access tool; SoftPerfect's Network Scanner is the most abused network enumeration tool, and WinRAR is the most abused archiving tool.

Microsoft binary abuse saw the least amount of difference year-on-year. Of the top 30 binaries observed, only two differed from last year, with the top 15 binaries appearing in nearly the same order. As with last year, we see a plurality of the top 30 tools (40%) being used for reconnaissance. The perennial leader, RDP, is still at the top, but we're happy to report that its overall use has declined. While still high, only 66% of cases involved internal use of RDP in attacks and 10% involved external use. We also saw a halving of exposed RDP systems. Chalk one up for the good guys, gals, and non-binary pals.

“Missing” logs — planting ice, harvesting wind

In the “everything else” (that is, all other findings) category, not much has changed either. The lack of MFA is still on top this year at 59% but has declined from last year’s 64%. Given that identity attacks are so prevalent, we need to keep driving this number down. All these complications seem to leave no choice.

Equally concerning are the cases where logs were missing, which was the second most common “other” finding. The leading cause of “logs missing” is still their general unavailability, but retention issues doubled over last year. This rise was largely driven by firewall appliances where the default was only seven days, and in some cases, 24 hours. While this may sometimes be due to hardware limitations, alternate storage should be arranged so they can be available for analysis when needed. And analysis needs to happen — you aren’t going to learn what you don’t want to know.

Unprotected systems came in at number three with 29.35% of cases. While we recognize that not all systems can be protected, most can. For those that can’t, additional mitigations and controls must be implemented.

This year also saw a tripling in the number of end-of-life systems implicated in attacks. This is another good example of basic hygiene issues that persist in many organizations. Unless end-of-life devices are absolutely necessary for the business, priority should be placed on reducing this kind of technical debt.

As the old saying goes, you don’t have to brush all your teeth, only the ones you want to keep. Taken together, these are unforced errors that must be remediated. We acknowledge that solving some of these problems can be difficult, but they are necessary if we want to increase resilience against attacks; it costs a lot to win, but even more to lose.

Artificial intelligence or virtual insanity?

Much ink has been spilled on the role that large language model (LLM)-supported generative artificial intelligence (GenAI) may or may not play in attacks in 2025. There were certainly some highly publicized research papers supporting adversary use of GenAI. Some turned out to be deeply flawed, while others misrepresented its importance in attacks, and some showed iterative development, but not revolutionary change — neither a rare nor different tune. While 2025 did see increased attacker use of GenAI on the social side of scams and attacks, it wasn't the watershed moment that signaled the era of the autonomous AI attacker. So far, enthusiasm is vastly outpacing evidence.

Anecdotally, phishing lures have been supercharged by GenAI. We can unfortunately no longer rely on poor grammar and spelling mistakes to detect a phishing email. To compound the problem, the combination of text and image generation has made brand impersonation trivial. Sprinkle in some highly personalized content scraped from breaches or publicly available information, and the ruse becomes flawless. Users must now rely mostly on intent and logic if they are to spot a phishing attempt. Generative AI also makes it so that the scale of these attacks can reach unprecedented levels. Deepfakes are now taking this even further; the quality of deepfakes is such that many victims would have a difficult time spotting one.

Proving its use is another matter. Outside of some deepfakes, it's nearly impossible to prove that GenAI was used to craft a phishing email. Even though it makes intuitive sense that attackers would attempt to make use of GenAI, there's no foolproof way to tell the difference between a meticulously handcrafted email versus one generated by GenAI. What GenAI does bring to the table is speed, volume, and democratization.

Using GenAI allows low-skilled attackers the ability to launch phishing campaigns at scale and faster than ever before. This also appears true for other aspects of the attack. Agents like HexStrike-AI and Cyberspike's Villager employ GenAI as an orchestrator to launch attacks using existing tools. As noted above, orchestration is a potentially rich direction for attacker to take GenAI; whether that's what's on tap for 2026… well, we'll know in a year.

According to Sophos incident responders, we saw a single verified and incontrovertible instance of GenAI use by an attacker — a video deepfake received via social media — in this year's case load. However, this case didn't result in an incident response investigation, because the victim promptly reported the interaction to the SOC. Score one for the humans.

While it seems inevitable that GenAI will someday cross the threshold into fully autonomous attacks, and possibly generate novel attack vectors and malware along the way, we aren't there yet. In the short term, the attacker gains will be — again — speed, volume, and democratization. There will be more malware — not necessarily better or more advanced malware, but likely orders of magnitude more of it. However, increased speed and volume, especially in the hands of unskilled attackers, can create noise as we have noted in the dwell-time discussion above. This creates an opportunity for defenders to brave the storm to come.

Does that meaningfully change how we defend our networks today? We don’t think so. The fundamental cybersecurity controls and mitigations need to be in place whether you are defending against a low-skilled cybercriminal, a nation-state adversary, or an AI-assisted attacker. Prevention through technology and policies plays a massive role. The required telemetry must be in place to spot the attacker, whether wetware-powered or software-powered. An understanding of what the telemetry is telling us is crucial. And the ability to act decisively on that telemetry is essential. We can run, but we can’t hide from it.

The obvious caveat is that the pace of change in this area means things can and will change without notice. But, barring advances in how GenAI fundamentally operates, these changes will be evolutionary, not revolutionary.

Case study: Time after time

Identification of phishing threats is often a core aspect of internal security training, with users getting wave after wave of “tests” designed to catch unwary clickers. It’s hard to claim it’s working, though — our statistics show that though phishing is still a relatively minor initial-

access vector as noted above, the year-to-year percentage of occurrences has doubled — from 2.13% in 2024 to 5.86% in 2025. (And that doesn't count all the cases that might have started with phishing but couldn't be confirmed and thus ended up in our "compromised credentials" catch-all.)

One case our MDR team handled last year showed just how much chaos phishing can bring to a bad situation. In that case, the victim's weak MFA configuration, coupled with a delay in credential resets, led to multiple waves of phishing attacks (and subsequent credential theft) during the same incident.

Fool me once

The first wave of phishing emails in this campaign was sent to multiple employees from an external Dropbox account masquerading as a trusted employee within the organization. The attacker leveraged the Dropbox "Share via email" feature to share a malicious PDF document across the organization. The name of the trusted employee in the subject line acted as bait — and multiple users bit, clicking on a malicious FlowerStorm URL contained within the document. Victims not only entered their credentials on the Dropbox portal but proceeded to accept an MFA prompt generated shortly after by the attacker.

This replay tactic is common among synchronous-replay adversary-in-the-middle phishing kits such as FlowerStorm. The counterfeit login page captures the victims' credentials and MFA tokens and then quickly replays these from the attackers' back-end server. The quick replay — often occurring within seconds of the legitimate request — tricks the user into accepting the attacker-generated MFA prompt.

Hiding in plain sight

Following the first wave of phishing emails, the attacker leveraged one of the newly compromised M365 accounts, created the same malicious PDF document within that employee's internal SharePoint, and sent a second wave of phishing emails out from the user's email address. Not only was the legitimate source now acting as bait, but the attacker had embedded the compromised employee's username within the

FlowerStorm phishing URL, aiding in the facade of legitimacy and luring additional users to take the bait. Multiple users once again interacted with the FlowerStorm phishing URL, allowing the attacker to capture more victim credentials and MFA tokens, initiate multiple successful M365 sessions, and create inbox rules that forwarded specified emails to the victims' RSS folders in order to conceal activity and delay detection. (There's nothing special about that folder, it's just one that's likely to be empty, since relatively few users set up RSS feeds these days.)

SharePoint did what?!

For a few years, attackers have been abusing file hosting services to facilitate delivery of malicious documents. To bypass various forms of analysis, the shared files are configured to be accessible only by the targeted recipient, or set to "view-only" mode. The former requires that the user be signed in to the hosting service; the latter prevents access to the contents of the malicious document by

Won't get fooled again (?)

The adversary did not stop there. A persistent M365 session was maintained due to organizational delays in remediation, and nearly a week later, a *third* wave of phishing emails was sent across the organization from *another* compromised user's legitimate email address. The phishing URL embedded in the new malicious document now contained the username of the new pawn, again fooling multiple new users. The pattern was nearly identical except for one thing: This time, no MFA prompt was presented to users.

Where did it go? Due to infrequent MFA prompt requirements and long-lived session token lifetime (the Microsoft Entra ID default token lifetime is 90 days, which though not as unreasonable as Slack's ten years is still long), the initial MFA approval from nearly a week prior was still valid. The attacker successfully replayed the stolen, unbound session token, enabling access without triggering an additional MFA prompt. The use of unbound session tokens was yet another downfall. Alongside strong MFA enforcement, bound session tokens should

email security solutions. Once a user clicks on a link within the malicious document, they are presented with an authentication prompt. In either case, the goal is to collect credentials and authentication tokens in cases where phishing-resistant MFA is not configured.

be enforced in M365 where possible, via Token Protection, to prevent the replay of stolen session tokens from other devices or locations.

Despite MFA being in place, the combination of weak MFA configuration, excessive session-policy timeout duration, and misjudged corporate priorities around credential resets played straight into the attacker's hands, allowing the adversary to persist in the network. Fortunately, the attack chain was stopped early by MDR response actions, leaving the organization nearly unscathed, though over one-fifth of the company's employees were touched in some fashion by at least one wave of the attack.

You can't always get what you want

There are some important takeaways here. Not only is it crucial to ensure MFA is *enabled*, but it is equally crucial to ensure MFA is strongly and effectively *configured* — for example, FIDO-based or other passwordless authentication, as well as smart choices such as frequent re-authentication. All it takes is one user missing from MFA setup or poor MFA prompt frequency to render the MFA trivial in an adversary's eyes.

With MFA determined to be either not enabled or not fully configured in 59.46% of incidents in 2025, this begs the question: Should education on MFA not also be an integral aspect of internal security training? This high unavailability percentage tended to encompass organizations in one of three categories: (1) organizations that believed MFA was enabled, but it was not at all; (2) organizations that had MFA enabled, but it was misconfigured; and (3) organizations who were aware MFA was not enabled. (The third gave investigators an array of reasons including management being "uncomfortable" with the concept.) MFA is a fundamental barrier against adversary initial access, with a decades-

long history and multiple means of implementation. Not allowing your organization to fall into any of these three categories because of “unavailability” is a crucial win.

Not all organizations are so lucky as to be left unscathed following the abuse of weak MFA, so if anything, take this as a sign — love your MFA, and lock it down.

And in the end

As we conclude another report, it's clear that defenders still have work to do: While there were some wins for defenders in 2025, you won't be finding us in the club, with bottles full of bub'. The lack of meaningful change in most of the year-over-year data means the attackers weren’t forced to try harder — they stuck with their same old tricks with little or no decline in success. Why go to the trouble of leveraging AI, really, when so many tried-and-true attacker TTPs are still working well?

Ironically, the abuse of identity is increasing because defenders *have* gotten better at patching some things. Attackers will always exploit our greatest weaknesses, but they remain flexible when times change. We defenders need to make it harder for them to operate — and to note where prevention can be reasonably applied to avoid having anything to detect (and to avoid appearing in the dataset of a future Active Adversary Report).

We can do hard things. Exploiting Flash and Java are a thing of the past. Thanks to Edward Snowden, most of the web is now encrypted. Those wins didn't happen overnight, but they happened. We can make things harder for the attackers if we systematically prioritize fixing the things that are under our control. And when you've done that, you go back, Jack, and do it again.

Acknowledgements

The authors wish to thank their colleagues Anthony Bradshaw, Ben Drysdale, Chester Wisniewski, Hamish Maguire, Jon Munshaw, Joshua Rawles, Karla Soler, and the analysts of the IR and MDR incident response teams.

Appendix: Demographics and methodology

For this report, we focused on 661 cases that could be meaningfully parsed for information on the state of the adversary landscape between Nov. 1, 2024 and Oct. 31, 2025. 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. For year-over-year data, we rely on AAR-gathered historical data going back to the 2020-era cases analyzed in our 2021 Active Adversary Playbook report.

Methodology

The data in this report was captured over the course of individual investigations undertaken by Sophos' X-Ops Incident Response and MDR teams. This includes cases handled by the Secureworks Incident Response team, including cases from the 3.5-month period before Sophos' acquisition of Secureworks in February 2025; those are aggregated with the rest of the IR cases (rather than MDR cases).

For this first report of 2026, we gathered case information on all investigations undertaken by the teams throughout 2025 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 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 661 cases, and those are the foundation of the report. The data offered in the downloadable dataset has been further redacted to ensure customer confidentiality.

The following 70 nations and other locations are represented in the data:

Angola	Cyprus	Korea	Singapore
Argentina	Dominican Republic	Kuwait	Slovenia
Aruba	Ecuador	Luxembourg	Somalia
Australia	Egypt	Malaysia	South Africa
Austria	Finland	Mexico	Spain
Bahamas	France	Morocco	Sweden
Bahrain	Germany	Netherlands	Switzerland
Barbados	Guatemala	New Zealand	Taiwan
Belgium	Honduras	Nigeria	Thailand
Bolivia	Hong Kong	Panama	Trinidad and Tobago
Botswana	India	Papua New Guinea	Turkey
Brazil	Indonesia	Peru	Turks and Caicos Islands
Canada	Ireland	Philippines	United Arab Emirates
Chile	Israel	Poland	United Kingdom
Colombia	Italy	Portugal	United States of America
Costa Rica	Jamaica	Qatar	Vietnam
Côte d'Ivoire	Japan	Romania	
Croatia	Kenya	Saudi Arabia	

Industries

The following 34 industries are represented in the data:

Advertising	Entertainment	Logistics	Rental & Leasing
Agriculture	Financial	Manufacturing	Retail
Communication	Food	Media	Services
Construction	Government	Mining	Sports
Consulting	Healthcare	MSP/Hosting	Telecom
Education	Hospitality	News Media	Transportation
Electronics	Information Technology	Non-profit	Travel and tourism
Energy	Insurance	Pharmaceutical	
Engineering	Legal	Real estate	

About the authors

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.

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.

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.