2026 Data Breach Investigations Report

verizon business

About the cover

“The only constant is change” is an aphorism commonly ascribed to Greek philosopher Heraclitus. There has been no historical evidence uncovered that he had any hands-on experience with cybersecurity, but he would be right at home in our field with this mentality. But even as the threat landscape constantly evolves and changes, the 2026 edition of the Data Breach Investigations Report (DBIR) invites you to consider the importance of the fundamentals of cybersecurity as the best way to brave all of this change. A little cyber-stoicism, if you will.

On our cover, you can see concentric rings, each one representing a year of our data, floating down and settling onto the foundation of our cybersecurity knowledge. They add to our understanding and complement our defensive strategies and are segmented by the incident patterns from the past four years.

Our own 2026 report is the topmost ring, followed by 2025, 2024 and 2023, the last one already settled into the foundation.

There are more zero days and critical vulnerabilities year over year (YoY), generative artificial intelligence (GenAI) augmented malware is now a common occurrence, and complex forms of social engineering are becoming more successful as the prelude to a breach. Their speed may be increasing, their scale might be a concern, but those are all challenges defenders have been facing for a long time. This new world should require more focus, more agility, but does not necessitate an upheaval. Refinement, not revolution. We will be ready for the future if we continue to collaborate and work together for the greater good.

Also, yes, those are technically donut charts. Sorry, not sorry.

2026

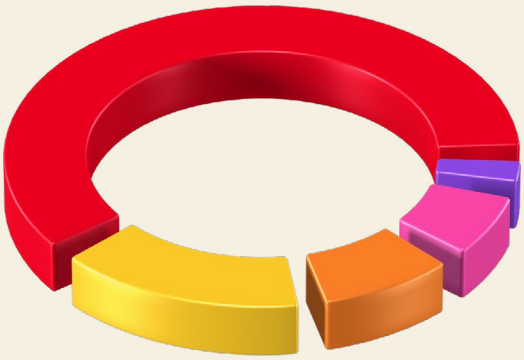

- **61%** System Intrusion
- **17%** Social Engineering
- **10%** Basic Web Application Attacks
- **8%** Miscellaneous Errors
- **3%** Privilege Misuse

2025

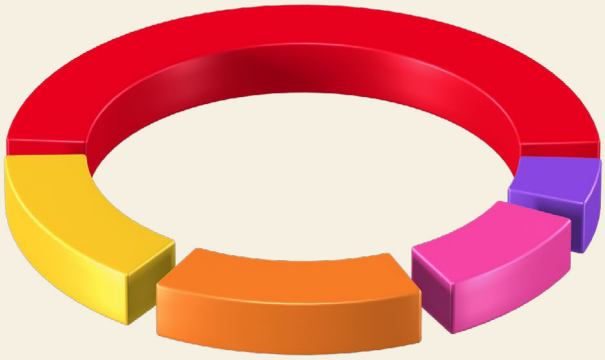

- **53%** System Intrusion
- **18%** Basic Web Application Attacks
- **17%** Social Engineering
- **12%** Miscellaneous Errors
- **7%** Privilege Misuse

2024

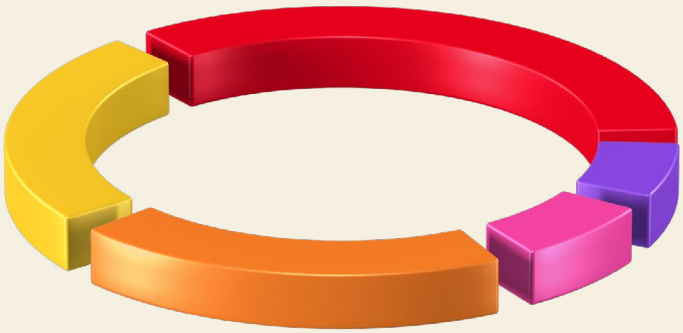

- **36%** System Intrusion
- **25%** Miscellaneous Errors
- **22%** Social Engineering
- **9%** Basic Web Application Attacks
- **8%** Privilege Misuse

Table of contents

Introduction

Welcome to Verizon's 2026 Data Breach Investigations Report! Hello again to those who've been with us over the years – and to those joining the DBIR community for the first time, it's great to have you. As always, we're glad you're here.

In this 19th edition of the Verizon DBIR, we dig into more than 31,000 actual real-world security incidents, of which more than 22,000 were confirmed data breaches involving organizations in 145 countries. This represents the largest number of breaches we have ever examined in a single report! Yes, we realize that we have said that before, but what can we say? It's still true because the number of cases we examine continues to increase YoY. We leave it up to you to determine if that is a good thing or a not so good thing.[1] For the victim organizations, it is undoubtedly the latter, but for our purposes of illuminating threats to your business, it is firmly in the former camp.

If we were to give this report an overarching theme, it would be "keeping a strong foundation in the face of change." Few people would argue that change, in every aspect of modern life, confronts us at an ever-increasing pace these days. The insights we try to provide in this report attempt to equip enterprises to meet cybersecurity changes in the most effective manner possible. And even though this report's dataset covers Oct 2024 through Nov 2025, both the DBIR team and Verizon are keenly aware of the growing impact and capabilities of AI-augmented vulnerability research and weaponization so far in 2026 based on early indicators and trends observed at the time of publication, and will provide some forward-looking commentary in regards to that where applicable.

We have observed that, in some areas, cybercrime has shifted in meaningful ways since the publication of the 2025 report. In others, it is less a matter of change and more a matter of speed and scale. Exploitation of vulnerabilities, discussed in several sections of the report, has now emerged as the most common way attackers gain initial access into an organization's environment, which underlines the ongoing importance of getting the basics right. Additionally, as the ancient prophecies[2] foretold, threat actors are increasingly relying on GenAI to assist them with various stages of their attacks, such as choosing targets, gaining a foothold within those targets, conducting vulnerability research, and developing malware and other tools to make their efforts more effective and efficient. Meanwhile, Social Engineering, a longtime fan favorite, is evolving, as well, with attackers increasingly using voice and other mobile-centric techniques to catch people off guard in the middle of the workday.

Regarding the System Intrusion pattern, we discuss the fact that Ransomware continues to be among the most disruptive and impactful types of breaches we see. Not unlike the price of everything from fast food to adult beverages in ballparks, it continues to trend upward. And we would certainly be remiss if we did not reference the increasing role that the broader web of third parties that organizations rely on can play in your security posture. However, there is a silver lining, but you will have to read on to learn about it. After all, we didn't spend all this time and effort writing jokes and witty content to simply give away the whole story on page one.

As we have done for the past several years, we examine key industry verticals in detail, along with a snapshot for small- and medium-sized businesses (SMBs). And last but certainly not least, we once again provide regional analysis for the Asia and the Pacific (APAC) and Europe, Middle East and Africa (EMEA) regions, so you can see how these trends show up in your own sector and part of the world.

Amid all this change, one message stays the same: The threat landscape will keep evolving, but the fundamentals still matter most. Organizations that stay grounded in strong cybersecurity basics (clear visibility into assets and third parties, disciplined patch management, and well-practiced response plans along with a culture that supports and enables secure behavior) are better positioned to handle today's realities and whatever comes next.

Sincerely,

The Verizon DBIR team
C. David Hylender, Philippe Langlois, Alex Pinto, Suzanne Widup

With special thanks to our Verizon colleagues:

- Chris Novak, for guidance and support all these years
- Steven Baskerville, Darrin Kimes and Jim Meehan from the Verizon Threat Research Advisory Center (VTRAC) team
- John Sandiford from Verizon Cybersecurity Architecture Australia

Additional recognition for some of our research partners this year:

- Raymond Carney and Scott Caveza from Tenable
- Saeed Abbasi from Qualys
- Jay Jacobs and Michael Roytman from Empirical Security
- Felipe Esposito and Alexandre Sieira from Tenchi Security
- Kyla Guru and Jacob Klein from Anthropic
- Simran Khalsa and Kelly Shortridge from Fastly
- Kellie Roessler, Michael Barnhart and Rajan Koo from DTEX

1. Or a fantastic thing! Ok, we're data breach geeks.
2. And by ancient, we mean predicted in the past two DBIR reports and mentioned a couple paragraphs ago.

How to use this report

First-time readers:

Before you get started on the 2026 DBIR, it might be a good idea to take a look at this section first. We have been doing this report for quite a while now, and we appreciate that the verbiage we use can be a bit obtuse at times. We use very deliberate naming conventions, terms and definitions and spend a lot of time making sure we are consistent throughout the report. Hopefully this section will help make all of those more familiar. If you are a longtime reader (thank you!) and are already familiar with how to use the DBIR, you are welcome to skip to the next section.

What you will find here

The Data Breach Investigations Report (DBIR) focuses on the analysis of anonymized cybersecurity incident data that Verizon collects every year from almost a hundred data contributors. Those data points are normalized using the Vocabulary for Event Recording and Incident Sharing (VERIS) framework (more about it on the right), which provides us a great foundation for statistical analysis of this type of data. Given the culture of secrecy (and just how difficult incident response is sometimes) that still permeates these cases, we often don't have all the very specific details of any given incident.

The breadth of data collection is what sets this report apart. Vendor-specific reports are able to talk very authoritatively and in great detail about the cases they investigated themselves, but here we are seeking to bridge different perspectives and contributor types – large incident response outfits, boutique forensics firms, law enforcement from local to country level, cyber insurance brokers and reinsurers – with the hope that it will get us closer to the capital T "Truth" of what is going on in the threat landscape. This poses unique challenges that we go over at length in our "Methodology" appendix, and sometimes in the content of the report itself.

Sections of the report

The report is divided into four large sections:

- In "Results and analysis," we will be focusing on the big picture of what happened in the previous year and exploring our complete dataset in each of the four main components of the VERIS framework (Actors, Actions, Assets and Attributes), with eventual guest appearances from other VERIS enumerations as applicable. This section should be useful and provide actionable information for all our readers, regardless of their industry segments or regions of the world.
- In "Incident Classification Patterns," we subdivide our dataset into patterns, which are shorthand for specific, very common incident archetypes with illustrative names such as System Intrusion or Denial of Service (DoS). This section is specifically helpful if you are looking for a deeper dive into those categories of incidents and seeking additional research and remediation guidance.
- In "Deep-dive analysis," we highlight long-form research we have done for this year's report that didn't fit well in any other section. Expect cork boards, lots of red string and analysis combining all sorts of different datasets from our data contributors.
- In "Industries," "Focused analysis" and "Regions," we focus our view of the dataset across different industry verticals and regions of the world and provide additional analysis on SMBs. These sections provide more specific analysis for the segment and should help folks in each segment to focus on where they might want to prioritize their efforts.

VERIS framework resources

The terms "threat actions," "threat actors" and "varieties" will be referenced often. These are part of the VERIS, a framework designed to allow for the consistent, unequivocal collection of security incident details. Here is how they should be interpreted:

Threat actor: Who is behind the event? This could be the external "bad guy" who launches a phishing campaign or an employee who leaves sensitive documents in their seat back pocket.

Threat action: What tactics (actions) were used to affect an asset? VERIS uses seven primary categories of threat actions: Malware, Hacking, Social, Misuse, Physical, Error and Environmental. Examples at a high level are hacking a server, installing malware or influencing human behavior through a social attack.

Variety: More specific enumerations of higher-level categories – e.g., classifying the external "bad guy" as an organized criminal group or recording a hacking action as SQL injection or brute force.

There are also "vectors" and "motives" and "categories," but we do our best in each section to ease folks into the nomenclature and try to make it clear how to interpret those terms. Also, any weird capitalization issues you may find throughout the report are referring to VERIS "Proper Nouns" and have specific meaning tied to them in the framework. As much as in the Fae world, true names have power here.

Learn more here:

- github.com/vz-risk/veris – features the framework's JavaScript Object Notation (JSON) schema with some usage, utility scripts, enumeration listings, mappings to Center for Internet Security (CIS) Critical Security Controls, MITRE ATT&CK and a VERIS Style Guide
- verisframework.org – a slightly more user-friendly website providing information on the framework with examples and enumeration listings

Incident vs. breach

We talk a lot about incidents and breaches and we use the following definitions:

Incident: A security event that compromises the integrity, confidentiality or availability of an information asset.

Breach: An incident that results in the confirmed disclosure – not just potential exposure – of data to an unauthorized party. A distributed DoS (DDoS) attack, for instance, is most often an incident rather than a breach since data is rarely exfiltrated. However, we realize that doesn't make it any less serious.

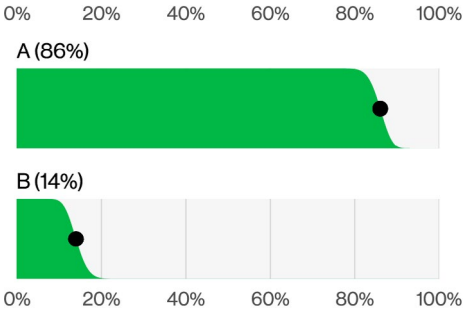

Figure 2. Example slanted bar chart (n=230)

Industry labels

We align with the North American Industry Classification System (NAICS) standard to categorize the victim organizations in our corpus. The standard uses two- to six-digit codes to classify businesses and organizations. Our analysis is typically done at the two-digit level, and we will specify NAICS codes along with an industry label. For example, a chart with a label of Financial (52) is not indicative of 52 as a value. "52" is the NAICS code for the Financial and Insurance sector. The overall label of "Financial" is used for brevity within the figures. Detailed information on the codes and the classification system are available here: census.gov/naics.

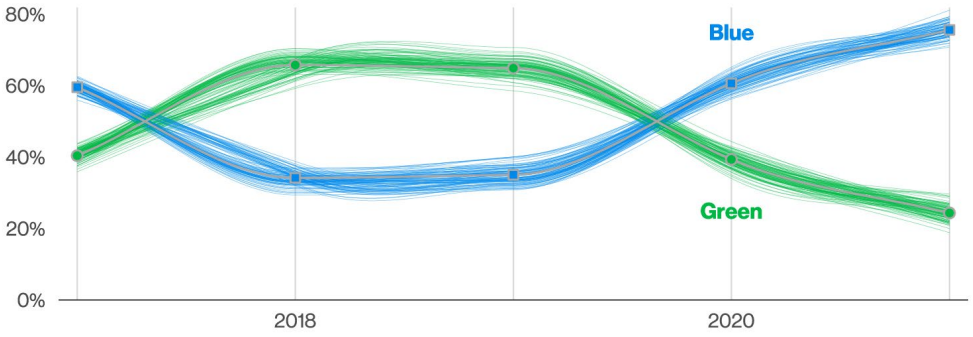

Figure 1. Example spaghetti chart

Being confident in our data

Starting in 2019 with slanted bar charts, the DBIR has tried to make the point that the only certain thing about information security is that nothing is certain. Even with all the data we have, we'll never know anything with absolute certainty. However, instead of throwing our hands up and complaining that it is impossible to measure anything in a data-poor environment or, worse yet, just plain making stuff up, we get to work. This year, you'll continue to see the team representing uncertainty throughout the report figures.

The examples shown in Figures 1, 2, 3 and 4 all convey a range of realities that could credibly be true. Whether it be the slant of the bar chart, the threads of the spaghetti chart, the dots of the dot plot or the colors of the pictogram plot, all convey the uncertainty of the cybersecurity industry in their own special way.

The slanted bar chart will be familiar to returning readers. The slant on the bar chart represents the uncertainty of that data point to a 95% confidence level (which is a common standard for statistical testing). In layman's terms, if the slanted areas of two (or more) bars overlap, you can't really say one is bigger than the other without angering the math gods.

Much like the slanted bar chart, the spaghetti chart represents the same concept: the possible values that exist within the confidence interval. However, it's slightly more involved because we have the added element of time. The individual threads represent a sample of all possible connections between the points that exist within each observation's confidence interval.

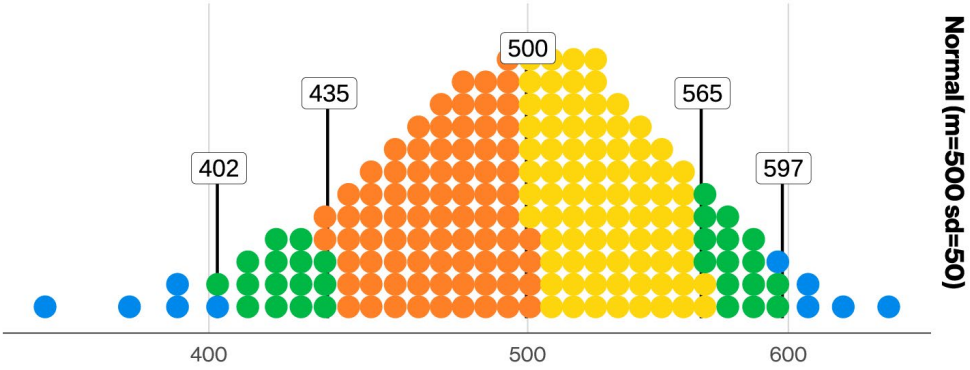

Figure 3. Example dot plot (n=10,000 – each dot is one event)
Orange: lower half of 80%; Yellow: upper half of 80%; Green: 80%–95%; Blue: Outliers, 95% of events: 402–597 80% of events: 435–565, Median: 500

As you can see, some of the threads are looser than others, indicating a wider confidence interval and a smaller sample size.

The dot plot is another returning champion, and the trick to understanding this chart is to remember that the dots represent a specific number of events, described in the figure caption. This is a much better way of understanding how something is distributed among organizations and provides considerably more information than an average or a median. We added more colors and callouts to those in an attempt to make them even more informative. In statistical terms, it's just a quantized density chart. In non-statistical terms, who doesn't love colored little dots?

The pictogram plot attempts to capture uncertainty in a similar way to slanted bar charts but is more suited for a single value or two. We hope they make your journey through this complex dataset even smoother than previous years.

This is what you could be looking at instead of unreadable pie charts everywhere else. Embrace the silly glyphs and never forget what they took from you.

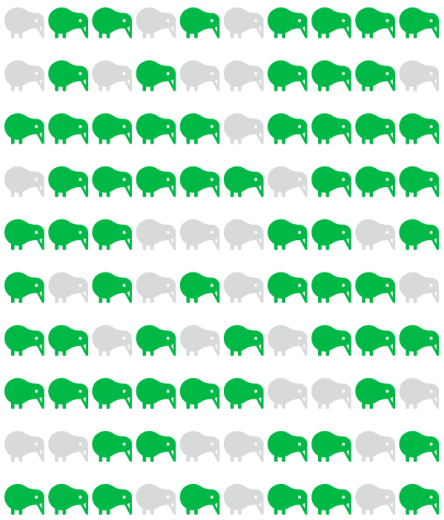

Figure 4. Example pictogram plot (n=100 – each glyph is one kiwi bird)

About the 2026 DBIR incident dataset

Each year, the DBIR timeline for in-scope incidents is from Nov 1 of one calendar year through Oct 31 of the next calendar year. Thus, the incidents described in this report took place between Nov 1, 2024, and Oct 31, 2025. The 2025 caseload is the primary analytical focus of the 2026 report, but the entire range of data is referenced throughout, notably in trending graphs. The time between the latter date and the date of publication for this report is spent in acquiring the data from our global contributors, anonymizing and aggregating that data, analyzing the dataset, and finally creating the graphics and writing the report. The jokes, sadly, do not write themselves.

Credit where credit is due

Turns out folks enjoy citing the report, and we often get asked how to go about doing it.

You are permitted to include statistics, figures and other information from the report, provided that (a) you cite the source as "Verizon 2026 Data Breach Investigations Report" and (b) the content is not modified in any way.

Exact quotes are permitted, but paraphrasing requires review. If you would like to provide people a copy of the report, we ask that you provide them a link to verizon.com/dbir rather than the PDF. You are, however, forbidden to generate pie charts based on data from the report. No exceptions.

Questions? Comments? Concerns?

Let us know! Send us a note at dbir@verizon.com or reach out to Verizon Business (or one of the authors) on LinkedIn. Be sure to tell your colleagues, families and neighbors (and Verizon Executives) about how much you love the report!

If your organization aggregates incident or security data and you're interested in becoming a data contributor or research partner to the annual Verizon DBIR (and we hope you are), the process is very easy and straightforward. Please email us at dbircontributor@verizon.com so we can discuss the details and make you a part of the DBIR research community.

Key topics and findings

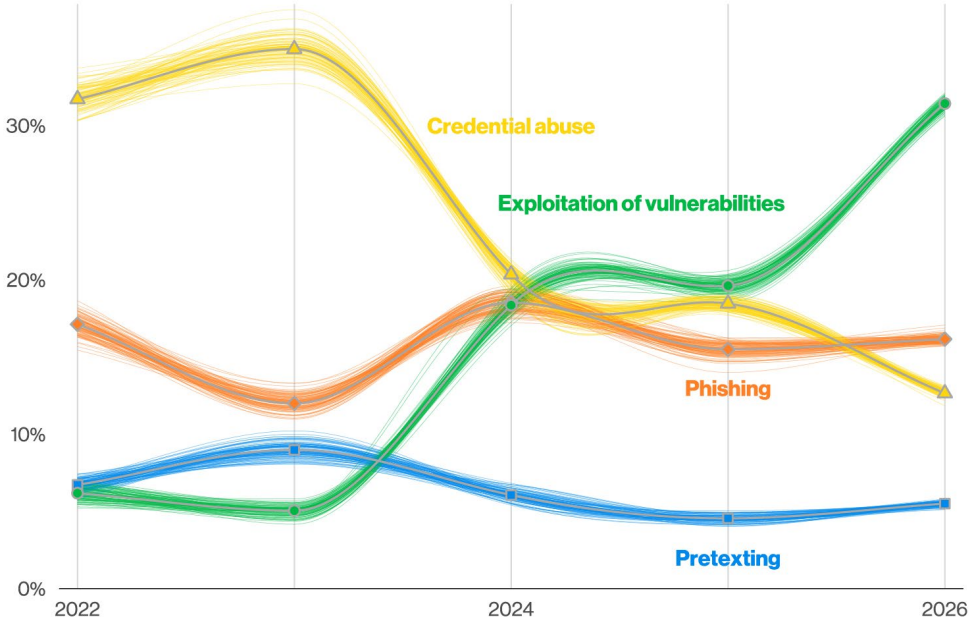

Figure 5. Known initial access vectors in non-Error, non-Misuse breaches over time (n for 2026 dataset=19,905)

Rise of vulnerability exploitation

Exploitation of vulnerabilities is now the most common initial access vector for breaches. It has risen to 31% in this year's reporting dataset, while credential abuse – the previous leader – is down to 13%.

Only 26% of critical vulnerabilities – defined as being in the Cybersecurity Infrastructure and Security Agency Known Exploited Vulnerabilities (CISA KEV) catalog – were fully remediated by organizations in 2025, a drop from the previous year's 38%.

The median time for full resolution went up to 43 days, almost two weeks more than the previous year's 32 days. In the median case, organizations had 50% more critical vulnerabilities to patch in this year's reporting dataset compared to the previous year.

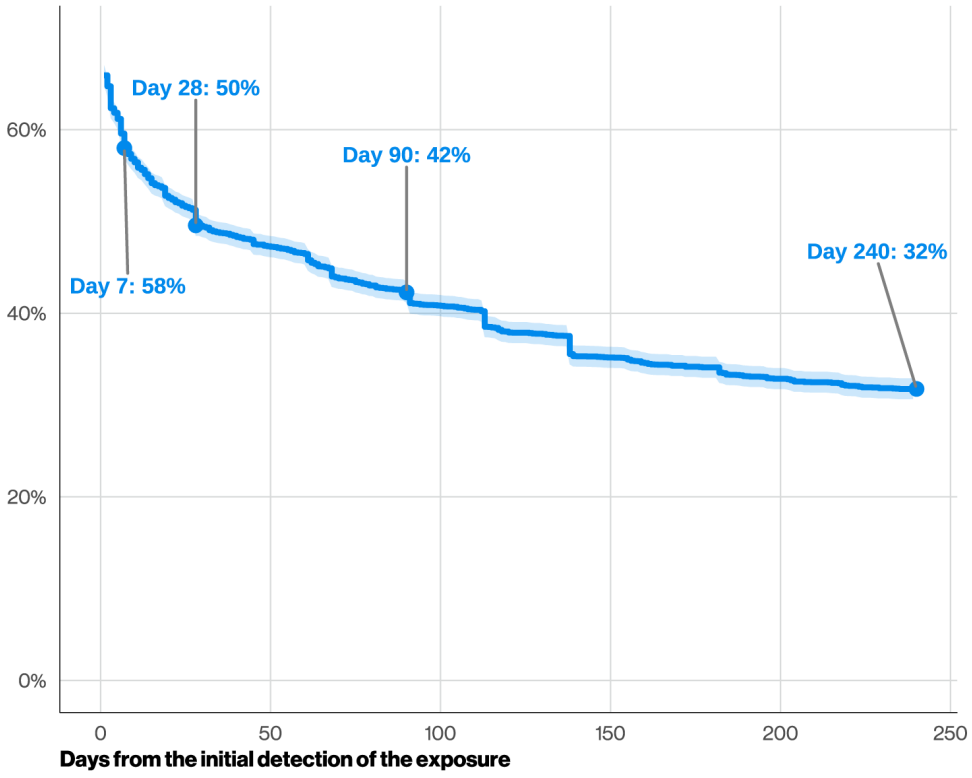

Figure 6. Survival analysis of third-party, cloud-based MFA exposures (n=7,513)

Growth in ransomware and third-party breaches continues.

Ransomware grew again to 48% of all breaches, up from 44% from the previous year. However, ransom payments have continued to decline among our dataset, as 69% of ransomware victims didn't pay. The median amount of ransom paid also continues a downward trend: $139,875 in this year's reporting dataset from $150,000 in the previous year.

As organizations increase their reliance on third parties for services and software, their exposure increases, as well, and breaches with third-party involvement have increased by 60% from last year's dataset, reaching 48% of total breaches.

Looking at remediation over time in third-party cloud exposure, only 23% of third-party organizations fully remediated missing or improperly secured multifactor authentication (MFA) on their cloud accounts, with 50% of all findings being resolved within a month.

For weak passwords and permission misconfigurations, the time to resolve 50% of all findings was much worse, reaching almost eight months.

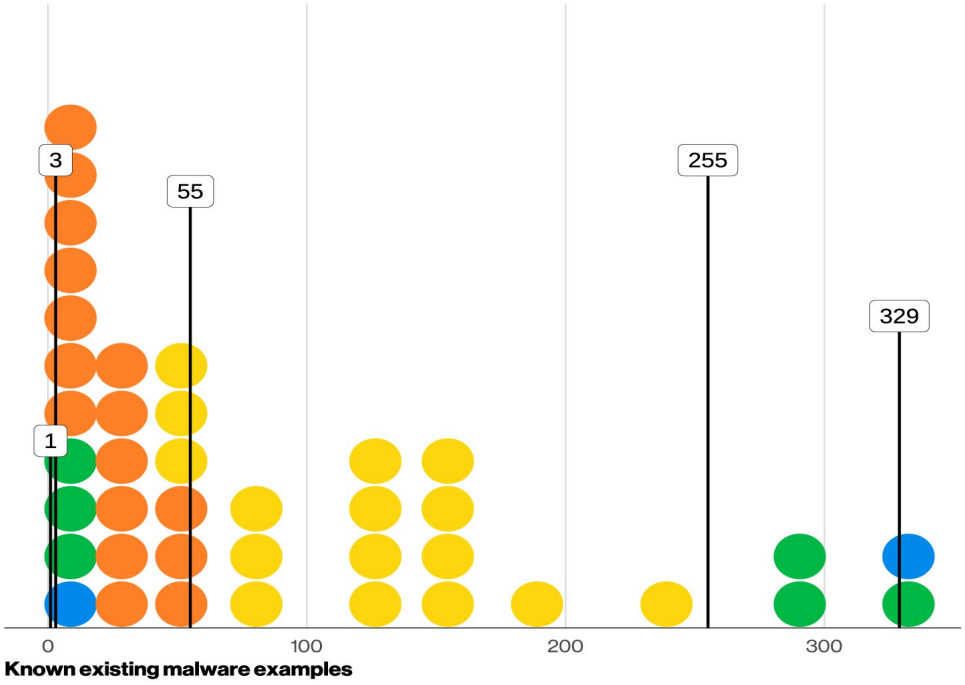

Figure 7. Distribution of known existing malware examples per ATT&CK technique observed (n=9,897—each dot is 247.43 observations)

Generative AI impacting the threat landscape

Threat actors are demonstrably using GenAI to help at different stages of attack, including targeting, initial access, and development of malware and other tools. The median threat actor researched or used AI assistance in 15 different documented techniques, with some Actors leveraging as many as 40 or 50.

Most AI-assisted development of malware and tooling was associated with well-known and defined attack techniques, with a median of 55 existing known malware examples performing the same functions.

Less than 2.5% of the AI-assisted malware observations involved less-common techniques with one or fewer known malware examples.

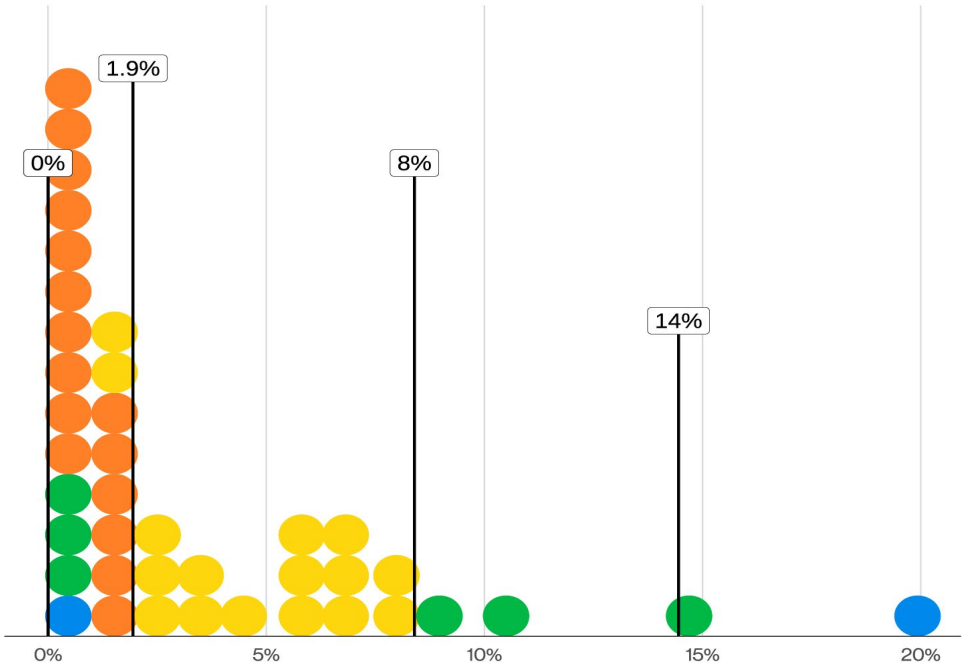

Figure 8. Distribution of success rate of non-Email vector-simulated social attack campaigns (n=35—each dot is 0.88 campaigns)

Mobile-centric Social Engineering

Human element was present in 62% of breaches, a slight increase from the previous year's 60%. Social Engineering was our third most common breach pattern, representing 16% of all breaches.

In phishing simulations, the median rate of successful "click" rates in mobile-centric vectors (such as voice and text messaging) is 40% higher than via email.

Pretexting has become a more common initial access vector to ransomware and extortion attacks. In all breaches, it reached 6%, while Phishing remained at 16% like the previous year. Pretexting is an attacker tactic in which a trusted relationship is built through concocted scenarios to trick the user into taking an action that unknowingly compromises the organization, frequently by voice communications but also seen via email or text messaging.

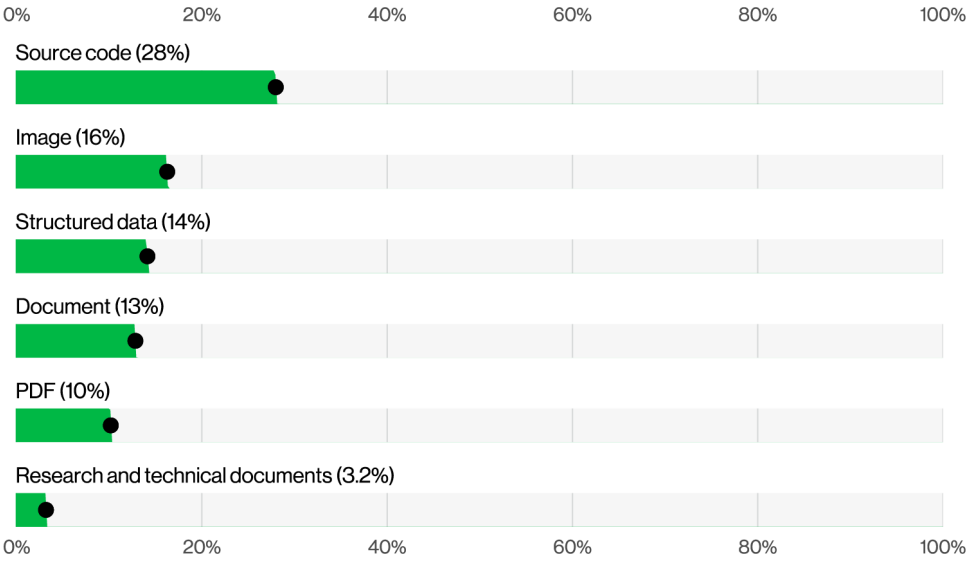

Figure 9. Select data types in untrusted DLP events targeting generative AI tools (n=858,440)

Shadow AI policy violations and malicious insiders

Regarding usage of unauthorized GenAI services ("Shadow AI"), 67% percent of users are using non-corporate accounts on their corporate devices to access AI services, a slight decrease from the previous year. However, 45% of employees are now considered regular users of AI (authorized or not) on their corporate devices, up from 15% in the previous year.

Shadow AI is now the third most common non-malicious insider action detected in our data loss prevention (DLP) dataset in 2025, a fourfold increase in percentage from the previous year.

The most common type submitted to external GenAI models was source code, followed by images and other types of structured data. In 3.2% of DLP policy violations, we even found research and technical documentation being uploaded to those unauthorized AI systems, which presents a risk of intellectual property exposure.

Results and analysis

/02

The big picture

Hello, everyone, and welcome to the "Results and analysis" section. This is where we cover the highlights we found in the incident dataset for this year's report. This dataset is collected from a variety of sources, including our own VTRAC investigations, incident reports and summaries provided by our data contributors and publicly disclosed security incidents.

Because data contributors sometimes come and go, one of our priorities is to make sure we can get broad representation on different types of security incidents and the countries where they occur. This ebb and flow of contributors obviously influences our dataset, and we will do our best to provide context on those potential biases where applicable.

This year, we have managed to analyze more than 22,000 breaches—a significant increase from our previous reports. Not only have we expanded our contributor base, but we have also doubled down on our capabilities to collect data in bulk from major public extortion and Espionage-motivated campaigns, along with ransomware actor activity. As a result, the report now offers a more expansive view of the threat landscape, although this information overload made our data pipelines run considerably slower.[3]

This subsection focuses on broader, and hopefully more actionable, high-level findings that go beyond the traditional structure of the VERIS 4As (Actor, Action, Asset and Attribute) and builds on some of the key metrics we have been highlighting over the past few years.

Hey kids, no name-calling please.

Longtime readers are likely aware that the DBIR team has always taken the position that we will not "call out" specific cases in the report and will refrain from including anything that would allow for inferring non-publicly available victim information.

This is very much still the case; however, for large-scale, publicly disclosed campaigns that affected a high number of organizations, we refer to the campaigns by their most commonly used terminology in the report to avoid confusion. We also comment on high-profile individual breaches but only refer to their publicly available information. Even if we had non-public information about those, we would be unable to correlate it with our dataset due to its anonymization.

Vulnerable beginnings and social changes

As any writer knows, there are few things more oppressive than a blank page staring back at you from your computer screen, as if mocking your inability to put your scattered thoughts down on paper.

Well, not today, blank paper! There is a good place to start this edition of the DBIR, and that is at the literal beginning of the breaches we analyzed. We have been tracking the initial access vectors of breaches for a few years now, and the continual growth of exploitation of vulnerabilities since the 2024 DBIR had us wondering when it would find its way to being the top vector. There is no need to wonder anymore.

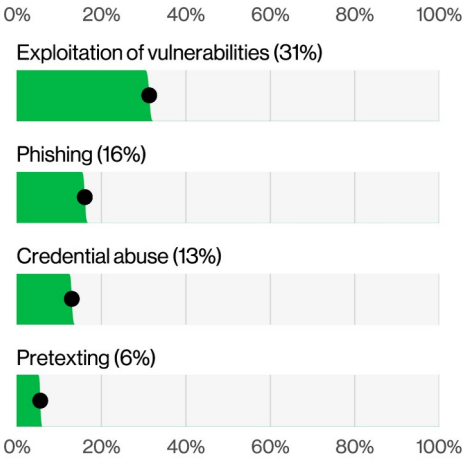

Figure 10. Initial access vectors—select enumerations in non-Error, non-Misuse breaches (n=20,023)

Figure 10 has all the details we need. The exploitation of vulnerabilities is the most prominent initial access vector in our dataset this year, reaching the height of 31%, up from 20% last year, which represents a 55% increase in this vector.

3. Have y'all seen how much computer memory and hard drives are costing these days?

Credential abuse, previously our most common, has fallen steeply from 22% in the 2025 DBIR to 13%, but there is a good contributing reason for that. This year, we added Pretexting – our second most frequent social action variety – to our tracked list of initial access vectors. Because there is frequently some overlap between Pretexting actions and credential abuse (this also happens in Phishing cases), this addition played a role in the lowered percentage for our former champion. For comparison with the 2025 DBIR results, this value without the addition of Pretexting would have been 16%.

However, this does not mean that defenders should discount the importance of mitigating credential abuse. This analysis focuses on the first initial action we can determine for the breaches we collect, but credential abuse is pervasive across various attack paths and is a legitimate mitigation target chokepoint. As Figure 11 demonstrates, if you consider all instances of credential abuse at any point in the breach progression, it still sits on top at 39%.

You should stop postponing that MFA rollout in your organization because credentials are an integral part of the threat actor's toolkit.[4]

It's not just a pretext.

The addition of Pretexting was not simply done on a whim.[5] Even though its total percentage is roughly the same as in the previous report (within statistical error), there was a significant number of high-profile ransomware breaches in this year's dataset that utilized this Social Engineering technique as the initial action.[6]

From this development, we believe it is important to re-introduce the distinction between the Phishing and Pretexting actions in the VERIS framework – and by extension in the DBIR – because the mitigations necessary for each of them might not be as similar as you think.

As VERIS defines it, Phishing is, in essence, an asynchronous social action. The victim will receive an email (or a text message) that attempts to alter their behavior in ways that will allow the incident to progress to its next step. "Enter your password here," "download this software there," whatever suits the threat actor's fancy.

Pretexting, however, is much more involved and insidious, as it intrinsically involves a synchronous component. There the actor is – on the other side of the phone, text message conversation or email thread – trying to convince the victim to do something or provide some information they shouldn't. The reason this matters is because the countermeasures and training needed to combat those two different scenarios are actually quite distinct.

Security awareness training involving email phishing simulation is quite ubiquitous in security programs nowadays, but the nature of pretexting requires more involved business-oriented rules and guidelines that are aligned with the nature of each potential target area in your organization. Training IT help desks and customer support agents to not be helpful and supportive in cases when a threat actor is trying to manipulate them is not as simple as "check if the email is external, from a source you trust and if it uses proper language."[7] We will be discussing voice and other non-email means of Social Engineering in the – drumroll – "Social Engineering" section of this report.

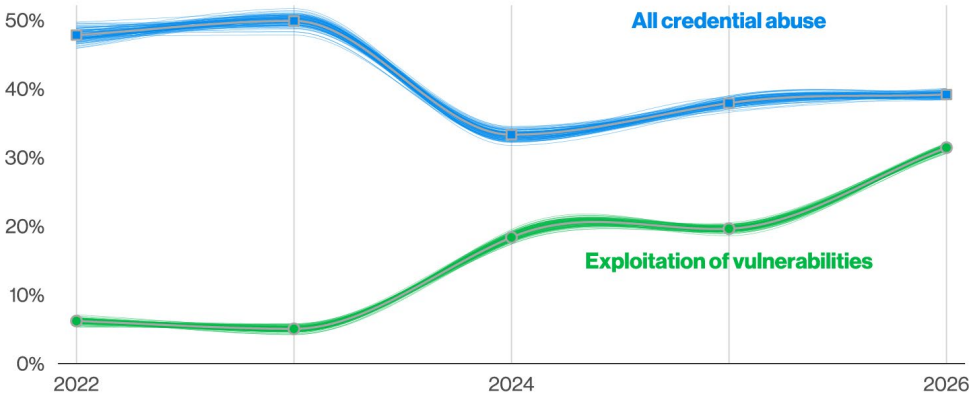

Figure 11. Select initial access vectors in non-Error, non-Misuse breaches over time (n for 2026 dataset=19,905)

Knowledge is half the "vulnerability" battle.

Given the uptick in Exploitation of vulnerabilities in our initial access vector analysis, it is a good idea to check in on our favorite sisyphean cause: vulnerability management.

To be clear, vulnerability management is an incredibly important risk mitigation process that needs to exist in virtually every organization, but the headwinds facing organizations implementing it have been discouraging, to say the least. Put quite simply, there are often too many vulnerabilities and not enough time for patching all of them.

4. If you are interested in learning more, we have some fascinating topics about infostealers and Initial Access Brokers (IABs) in the "System Intrusion" section.
5. There are several of those all the time in the DBIR, but this one is not one of them.
6. We will discuss them a bit further in this section.
7. Given the reported increase of AI assistance in phishing emails, we are changing the detection guidance from "does it contain many typos" to "does it contain em dashes."

We once again chose the CISA KEV list as our subset of vulnerabilities[8] that all organizations are incentivized to patch, and we set out to replicate the metrics of percentage of remediation and the median time to do so. In this study, which includes aggregated information from more than 13,000 organizations, we were not only focusing on vulnerabilities added to the CISA KEV in 2025. Any vulnerabilities uncovered by the scanners in an organization's environment that were in the CISA KEV by the end of 2025 were being counted. After all, an old unpatched vulnerability that suddenly becomes the focus of an attacker campaign can be as disruptive to the process as a brand new one.

Figure 12 shows the percentage of unique CISA KEV vulnerabilities found in organizations per remediation status.

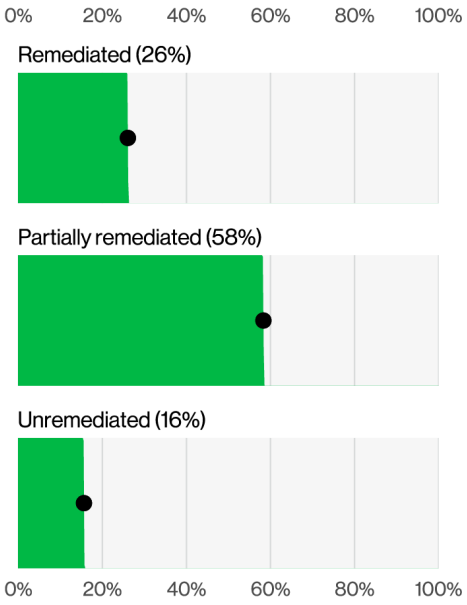

Figure 12. CISA KEVs per CVE resolution status (n=515,170)

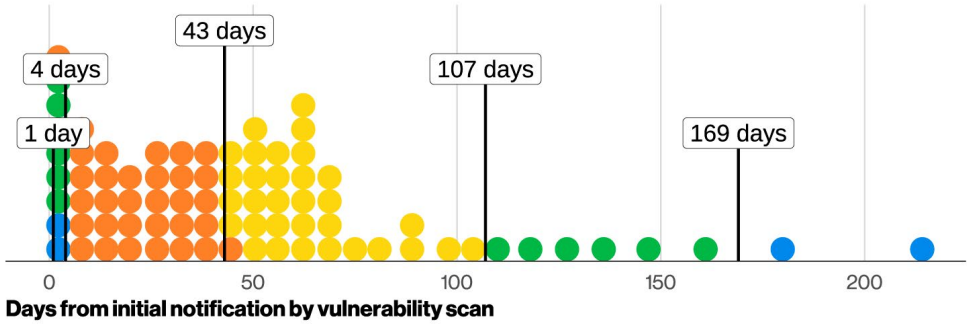

Figure 13. Distribution of the median of days until full remediation of CISA KEV vulnerabilities in a single company (n=10,597—each dot is 132.46 unique CVEs per company)

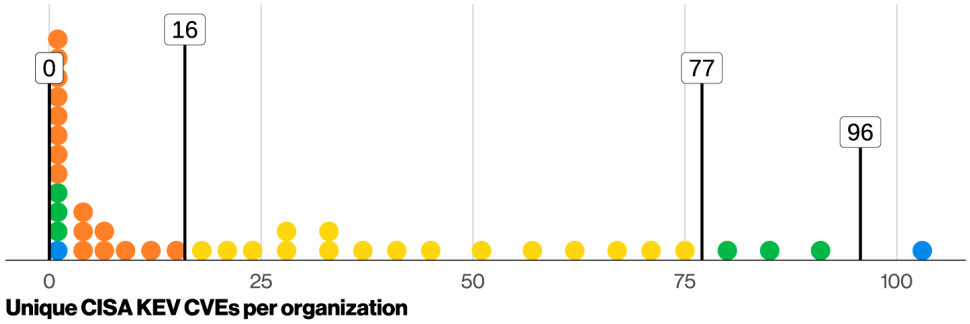

Figure 14. Distribution of unique CISA KEV CVEs per organization (n=13,773—each dot is 247.43 organizations)

The "Remediated" column is our favorite because it represents that the organizations in question have fully patched all instances of those specific vulnerabilities.

In the "Partially remediated" column, only some of the instances were patched for whatever reason, and maybe there were valid risk management-related reasons not to patch them all, as opposed to a simple failure to complete the task. Of the "Unremediated," we dare not speak. That way lies madness.

The results are worse than last year. Only 26% of the CISA KEV vulnerabilities had been fully remediated, a considerable drop from last year's 38%. Even if you are a "glass half full" person and want to give full marks to partially remediated ones, the unremediated ones add up to 16%, an increase from the 12% found last year.

There is also a worse result for the median time elapsed for a vulnerability to be fully patched by detection, shown in Figure 13. Our new median time is 43 days, almost two weeks longer than last year's 32 days. Figure 14, however, begins to elucidate the mystery: The median number of KEV vulnerabilities that had to be patched by organizations has risen in 2025 to 16, where this figure was 11 in 2024. That is almost 50% more KEV vulnerabilities to patch in a year.

8. To allow for easier comparison and aggregation of the multiple vulnerability management research partners the DBIR has, the singular focus was on vulnerabilities that have an assigned Common Vulnerabilities and Exposures (CVE) identifier. We know there are exploited vulnerabilities with no CVE assigned and also the larger discipline of exposure management, but we need to be able to compare apples to apples here across disparate datasets.

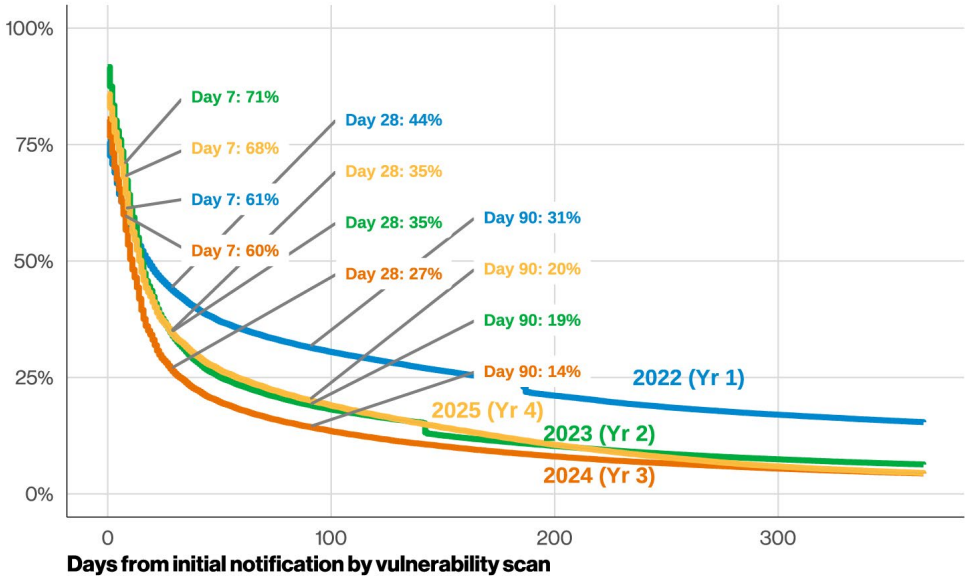

Figure 15. CISA KEV vulnerability survival analysis: four-year comparison (2022–2025) (n for 2022=68,697,749 vulnerability instances) (n for 2023=120,803,360 vulnerability instances) (n for 2024=295,795,092 vulnerability instances) (n for 2025=527,255,454 vulnerability instances)

Survival of the vulnerable

Those findings reinforce the "patching capacity issue" hypothesis we have proposed in our vulnerability management analysis section over the past few years. To shed more light on this subject, we decided to go back to a survival analysis approach to allow us to measure the full life cycle of each patched vulnerability instead of a snapshot at the end of the year.

Figure 15 shows the survival curve – the percentage of KEV vulnerabilities still open at weekly intervals – grouped by the last four DBIR reporting periods starting with the 2022 report. This dataset combines more than 1 billion anonymized vulnerability detection records, courtesy of one of our vulnerability management research partners.[9] The picture it paints is that of a treadmill picking up speed.

The 2025 DBIR, based primarily on data from 2024, was the high water mark. At every milestone in the survival chart, organizations were remediating faster than they ever had before, showing improvements from 2022 to 2023 and from 2023 to 2024.

Then 2025 happened and the curve shifted back to 2023 levels, with 35% still open at Day 28 (up from 27% in 2024) and with the long tail settling at 9%. This represents 47 million vulnerability instances that, based on our curve trajectory, are simply not being addressed any time soon.

In aggregate, the conclusion seems to be that organizations collectively did get worse at this patching thing, but what tipped the scales here was the volume flowing through the system.

There were 68.7 million records in the 2022 dataset and 527.3 million in 2025 – almost eight times the volume.

At Day 28, that 35% translates to 184 million open vulnerability instances, up from 31 million in 2022. The number of distinct organizations in this dataset did not vary significantly YoY, so we are comfortable in comparing those absolute vulnerability numbers to provide perspective on the enormity of the task they face.

Another interesting way to measure this is by focusing on the top performers. Organizations with well-developed vulnerability management processes do not typically wait for a vulnerability to appear in the CISA KEV catalog before committing to patching it.

In the 2024 calendar year, 17% of vulnerability instances were remediated prior to their inclusion in the KEV catalog. However, despite organizations' best efforts, this preemptive remediation rate fell to 12% in 2025. In absolute numbers, defenders proactively patched a staggering 63.7 million vulnerability instances in 2025, a 30% increase from 2024 (48.9 million).

There appears to be a ceiling, and the data suggests diminishing returns at current resource levels. By Day 7, which is an incredibly fast milestone for remediation by any standard, somewhere between 60% and 70% of KEV vulnerabilities remain open regardless of year, volume or organizational maturity. This first-week rate barely moved despite three years of additional process development, tooling investment and mandate pressure.[10]

Those results should inspire additional analysis by other researchers, but more than 1 billion records are nothing to sneeze at. This might be an initial measurement of the "speed of light" – a theoretical limit – for vulnerability remediation processes.[11] Organizations at their very best only get to fix 30%–40% of KEV instances in the first week after detection, so choosing[12] the correct ones to patch really is the key strategy.

9. Qualys published additional insights from this dataset in a report they published in March 2026: qualys.com/forms/whitepapers/the-broken-physics-of-remediation
10. Let's be optimists for a moment and assume organizations continually improve their processes.
11. Huge opportunity for a cybersecurity company to invent the Alcubierre drive: en.wikipedia.org/wiki/Alcubierre_drive
12. Or guessing

Recency bias for the win?

On the topic of strategies for choosing which vulnerabilities to patch first, the team was inspired by a recent report from one of our research partners[13] about the concept of resurgent vulnerabilities. The main concept is that even vulnerabilities that are very old – discovered years ago and that in an ideal world should have been patched already – will sometimes have exploitation activity associated with them all of a sudden.

At a high level, their results suggest that focusing only on patching the newer vulnerabilities and leaving a backlog of critical ones with no apparent exploitation is not a guaranteed win. However, this just gives us more tasks to complete, which does not help the issue of our patching needs outstripping our capacity to effectively patch. Is there a more principled way of deciding against new versus old?

To replicate this analysis with frequency of exploitation data from another of our new research partners, we enlisted one of the former DBIR writers who happens to work at said research partner.[14]

There were 1,526 CVEs listed on the CISA KEV as of February 2026, and 991 of those had some exploitation activity over the past 12 months.

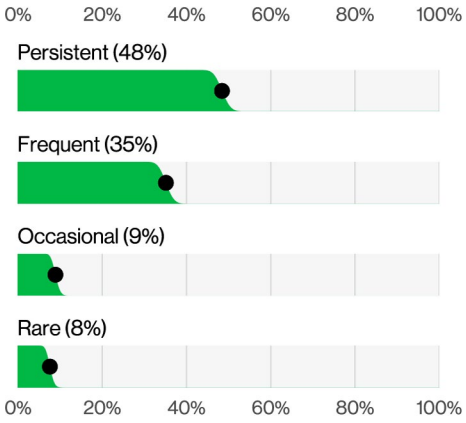

Figure 17. Clusters of expoitation activity frequency of CISA KEVs (n=991)

We will be focusing on those for the analysis, as the dataset did not (or could not) detect exploitation activity for the remaining vulnerabilities.

Figure 16 shows the frequency of exploitation in a very colorful chart, similar to the ones in the GreyNoise report we mentioned earlier. With the help of some old-fashioned unsupervised clustering algorithms, we can boil this down to four discrete frequency categories, described in Figure 17, for the CISA KEVs with detectable exploitation activity.

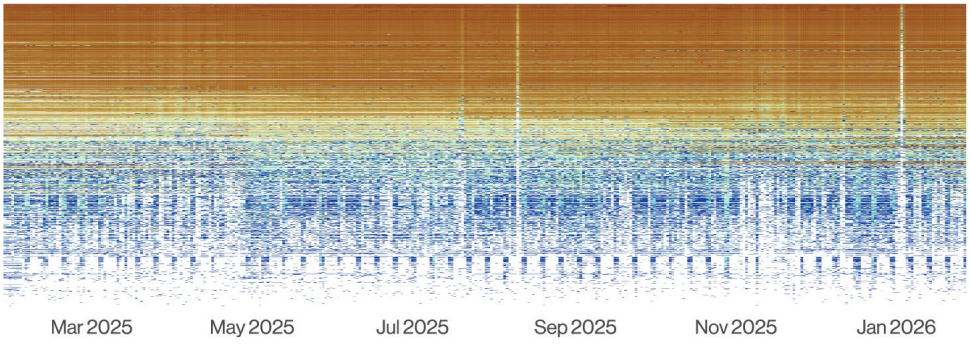

Figure 16. Heat map of exploitation activity frequency of CISA KEVs per day (n=991)

It should shock no one that nearly half of the vulnerabilities in the KEV are deemed to have "Persistent" exploitation, which means they could be detected an average of 96% of the days. What might prove surprising, though, is how many of those were older vulnerabilities and not just the latest and shiniest. Only 20% of the vulnerabilities in the "Persistent" category were registered in the CVE database in 2024 and 2025. For the other 80%, organizations would have had about two years' advance notice in which to patch.

The challenge is that the CISA KEV is a timestamp and not a timeline. It marks when a vulnerability reaches a critical mass of exploitation in the wild but does not discuss if the exploitation rate falls or disappears. Let's try to add this time dimension to it because, regardless of the categorical label on a vulnerability, the same decay pattern seems to hold once exploitation is observed.

Using the same frequency analysis of vulnerabilities as before, and some light modeling, it is possible to try to forecast how likely it is for a vulnerability to resurge and go back to being actively exploited based on how recently it has been exploited.

The analysis covers roughly 1.4 million observations of approximately 1,000 vulnerabilities over six years. For each day a vulnerability was being tracked, we recorded two things: how many days had passed since the last known exploitation activity (so we naturally had to wait for the first observed activity) and whether that vulnerability was exploited again within the next 30 days. The model shown in Figure 18 finds the mathematical relationship between those two things.

The main conclusion here is that the longer it's been since a vulnerability has been exploited, the less likely it is to be exploited again soon. A case of recency bias if we've ever seen one but surprisingly aligned to the measurements.

13. Read the GreyNoise report at greynoise.io/resources/how-resurgent-vulnerabilities-jeopardize-organizational-security.
14. Shout out to Jay Jacobs, Chief Data Scientist at Empirical Security. Much of the automation the DBIR team uses today leverages code he wrote more than a decade ago, and we will never forgive him for that.

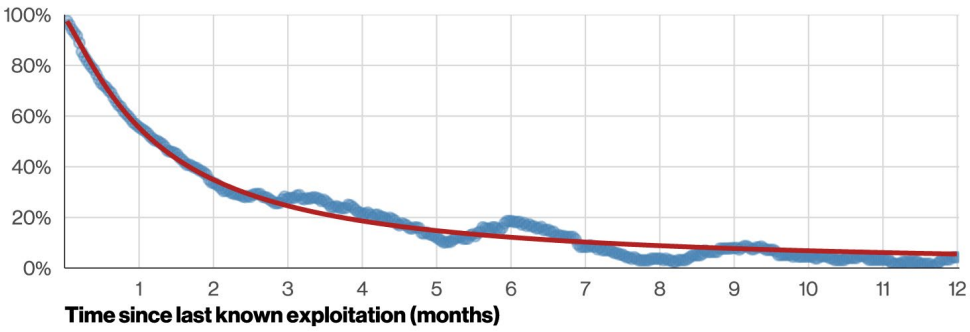

Figure 18. Re-exploitation probability by last known exploitation (points=empirical rate by observation count, line=GLM fit)

The pattern held consistently across the full six-year observation, which gives reasonable confidence it reflects an actual structural feature of exploitation behavior rather than just noise.[15]

The odds drop off quickly in this model. The probability drops by roughly half at 30 days, again at 90 days and again by half around nine months. After around a year, the probability of seeing resurgent exploitation activity is about the same as if it wasn't ever exploited.

With this result, if faced with the choice of patching a vulnerability that is less than a year old in the KEV but that hasn't been exploited recently or one that isn't on the KEV (yet) that your threat intelligence indicates does show recent exploitation history, focusing on the one with recent activity could be a smarter bet.

This approach is informed by historical patterns and should be weighed alongside other risk factors specific to your environment.[16] This model behavior does align with the understanding that threat actors have to develop exploits and maintain infrastructure scanning for vulnerabilities, and there are only so many of those that can be done at the same time.

However, this is a base assumption that the promise of increased automation of vulnerability discovery and exploit development from GenAI tooling could upend. It's a fun way to look at the problem, regardless.

Two's company, three's a breach.

In addition to our usual discussion of initial access vectors, there are a few other metrics we have been tracking as part of our big-picture analysis. We have a visual aid for you in Figure 19, where we capture major data points in our dataset, including the first item, which serves as an important reminder of the prevalence of credential abuse even in the face of growing vulnerability exploitation.

The second data point measures the non-intentional human element contribution in breaches, which has been fairly stable since the 2024 DBIR. This year it reached 62% of all breaches. While this is technically a statistical increase from the 60% recorded in the 2025 DBIR, when one accounts for our error margins, it definitely does not warrant extensive analysis.[17]

People are not computers. Plan accordingly in your cybersecurity strategy. And please feel free to refer back to our discussions of this topic in the 2024 and 2025 DBIRs.[18]

The third and last measure concerns the involvement of third parties in breaches, and we will dedicate time to discussing it in the next few pages. This metric comes in at 48% this time around, up from 30% last year. That is a 60% increase, after already doubling the year before – quite a trajectory.

This sustained growth has proven impossible to ignore, as many of the year's most high-profile and well-publicized breaches involved multiple third parties. In several of the more notable campaigns, attackers compromised more than one third-party provider at the same time.

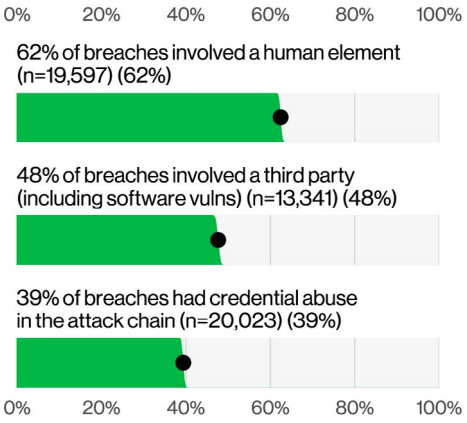

Figure 19. Select key enumerations in breaches

15. Or the recency bias devils trying to clean up their act
16. The insights offered by the DBIR should not be used for vulnerability betting in prediction markets.
17. If you're disappointed to hear this, don't fret. We have a whole section dedicated to Social Engineering.
18. Unfortunately, we have to adhere to a very strict page budget to get this report out the door each year.

The rule of three[19]

As a reminder, our third-party metric combines three different kinds of business relationship archetypes, which concern themselves with where the initial access happens and where the data that was breached was stored:

1. **Vendor in an organization's software supply chain:** The data and initial vector were under the organization's custody and control, but the initial vector was only made possible given a vulnerability in a vendor product (i.e., the majority of single Exploit vuln actions we mention in this report) or the compromise of said vendor and the inclusion of a back door in the software (such as what happened with SolarWinds a few years ago, based on publicly disclosed information).
2. **Vendor hosting an organization's data in its environment:** The initial access is against the vendor itself, and the vendor is the custodian of your organization's data. The two big sub-categories here involve your vendor's infrastructure being breached directly or your account to the vendor's environment being stolen and used against its systems (last year's campaign involving Snowflake is a good example of the latter, based on publicly disclosed information).
3. **Vendor with connection to an organization's environment:** The initial access is on the vendor, and lateral movement is done to reach the data inside your organization. This can either be a literal network connection that is leveraged by attackers (like the Target breach that happened more than a decade ago where a network connection to a vendor was the way in, based on publicly disclosed information) or by the vendor losing[20] credentials to an organization's internal systems to attackers.

The bad news is that we increasingly see a combination of two of those – or even all three – contributing to a breach. Based on publicly disclosed information, one such example involved a recent campaign around the Salesforce plugin Salesloft Drift facilitating breaches of customer data in the platform.[21]

According to publicly available information on this campaign, the customer OAuth tokens (or the keys to derive those OAuth tokens at will) from the Salesloft Drift application were compromised (Archetype 3, initial access vector against the Salesloft vendor) and then they were used against the Salesforce platform to steal data from the customers (Archetype 2, your data exfiltrated from the vendor environment).

Avoiding data breach spectator mode

At first glance, there doesn't appear to be anything that could have been done to prevent these from the victim organization's perspective.

But with a closer analysis of the root causes, a good number of these cloud-based, third-party incidents highlighted in the media in 2025 boil down to insecure authentication (absence of MFA, improper credential rotation) or lack of least privilege enforcement for users or service accounts.

Excessive privileges in cloud environments – be they Infrastructure, Platform or Software as a Service (IaaS, PaaS and SaaS, respectively) – is a pervasive issue. In fact, any considerations on authentication, secret management and obviously MFA are strong points of attention for any cloud environment. Since we are in the third-party section, let's dive into what can happen in your third-party's cloud environment configuration, aka, their third party.[22]

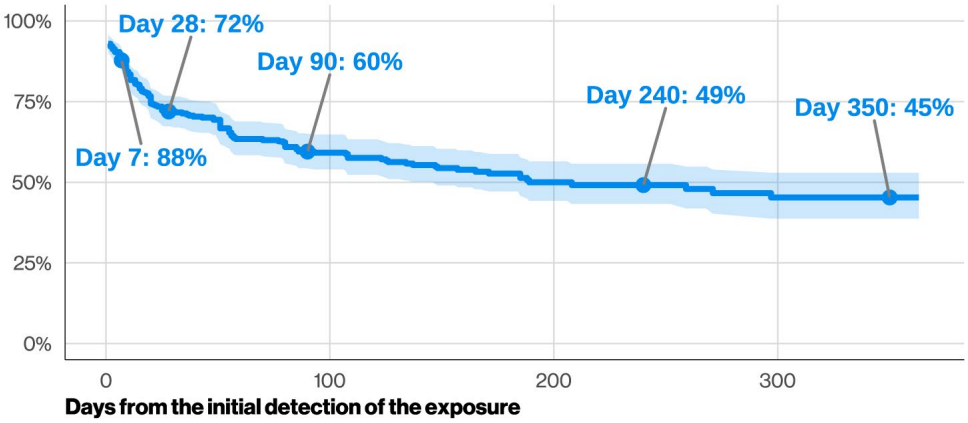

Figure 20. Survival analysis of third-party, cloud-based authentication and privilege exposures (n=354)

19. Not to be confused with The Sign of Four
20. Or handing over due to Social Engineering
21. help.salesforce.com/s/articleView?id=005134951&type=1
22. As our 2025 DBIR cover would like to remind you, it is third parties all the way down.

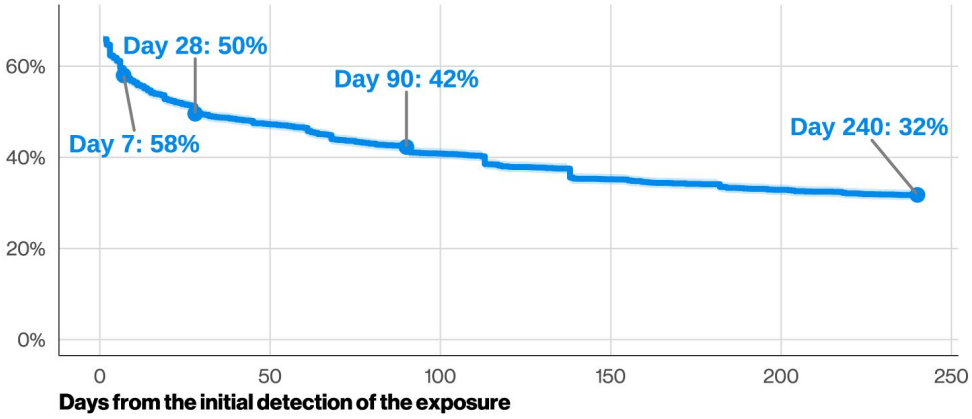

Figure 21. Survival analysis of third-party, cloud-based MFA exposures (n=7,513)

The third rule of survival

This year, we have access to a robust dataset from a new research partner focused on third-party cyber risk management (TPCRM), including resolution times of common types of exposures collected directly from inside the third-party cloud environments – a remarkable find because, up until now, we've only had access to outside scanning data of those kinds of environments.

Resolution times allow us to do survival analysis and thereby pull off the same tricks we usually have available for vulnerability management. Figure 20 includes the survival analysis of poor password practices and excessive access permissions combined on those third-party cloud environments. Taking almost eight months to fix this category of exposure makes our vulnerability management numbers seem lightning fast. Still, the full remediation percentage here is 31%, comparable with our success rate with the CISA KEV vulnerabilities in the initial access vector analysis discussed earlier in this section.

The lingering tail remaining so high – close to our 50% survival rate – shows that this is an exposure that organizations struggle with resolving. In fact, what is an "excessive permission" anyway? An admin account is an easy concept to understand, but given the granularity present in cloud environment configurations in IaaS environments,[23] it is an almost impossible task to make informed decisions based on the least-privilege principle under normal business time constraints.

Figure 21 shows that with something that is commonly accepted as a control that should be implemented at every opportunity, such as correct MFA configuration, the overall remediation results are better. Our 50% survival rate here is roughly a month, and the tail converges to roughly 32% of lingering MFA-related issues. The percentage of organizations that fully remediated this category of issues is a bit lower, sitting around 23%.

For a different perspective, looking at a point-in-time snapshot across another cloud exposure dataset, 37% of organizations had an admin account with MFA disabled on an IaaS offering. For reference, only 14% of organizations had an admin account with MFA disabled on Snowflake, leading us to believe most customers got the memo from last year's breach campaign.

This is just a small slice of potential exposures of cloud environments, but a strong starting point is to focus on the authentication and authorization layers, as those are usually the ones that end up on an organization's end of the responsibility matrix of cloud environments. These are security fundamentals that have been understood and had measurable success for decades now,[24] and they should be applied to all systems and environments that allow them, where feasible and appropriate to the organization's risk profile. We should pay special attention to service and machine accounts, as those will likely be the ones leveraged in our potential agentic AI future.

23. Not to mention the incredibly complex user interface and experience (UI/UX) involved in taming access control in those cloud environments.
24. At least our Certified Information Systems Security Professional (CISSP) Exam Prep book from 25 years ago said this, so we can confidently say decades, plural (RIP, Shon Harris).

VERIS Actors

Throughout this report, we use terminology from the VERIS framework.[25] At the highest level of VERIS categories are the 4As: Actor, Action, Asset and Attribute. When we say "Actor," we mean exactly what it sounds like – the threat actor behind the breach or incident. Those actors fall into a few familiar buckets: External (someone outside the organization), Internal (an employee or other insider), Partner (an entity with a business relationship with the victim) and, finally, multiple (any combination of the above).

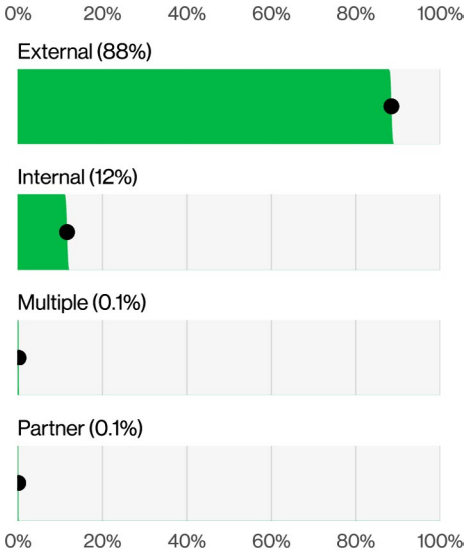

Figure 22. Actors in breaches (n=22,345)

You don't need a calculator to figure out that, no matter how big your organization is, there are always more people on the outside than the inside. Therefore, it shouldn't shock anyone unduly that External actors continue to be the ones causing the most trouble (Figure 22).

That's been the story for all 19 iterations of this report, and this year is no exception: 88% of threat actors were External – a bit higher than last year but still slightly below many of the earlier years when that number routinely sailed past 90%.

Internal actors appeared in 12% of breaches this year – a decline from last year's 18% but still enough to keep IT and security teams from chalking it up as a rounding error. These cases can be especially thorny, since they typically involve people with legitimate access and at least a basic understanding of how things work on the inside. By contrast, breaches involving Partners or multiple actors were relatively uncommon and barely moved the needle when compared to the dominant Internal vs. External storyline.

Who would do such a thing?

Each of the top-level Actor categories can be sliced into more specific varieties. For external actors, most roads lead back to Organized criminal groups (Figure 23). In this context, Organized crime isn't about cinematic mob families; it simply means criminals who run their operations in a systematic, repeatable way. They have a process, and they follow it with great success. This variety dominates the dataset, driven largely by the continued popularity of ransomware and other extortion-centric attacks.

State-affiliated actors make up the bulk of the remainder, appearing in close to 15% of breaches.

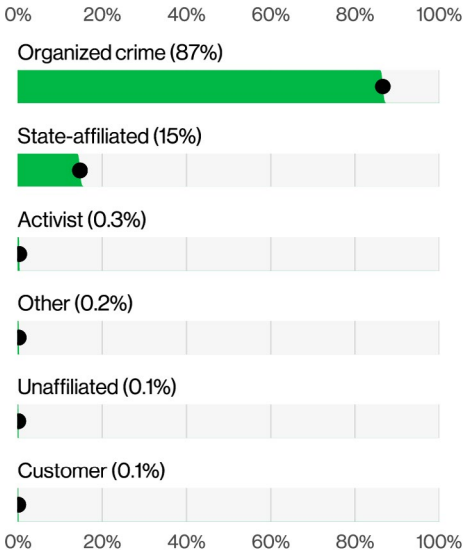

Figure 23. Top External actor varieties in breaches (n=16,491)

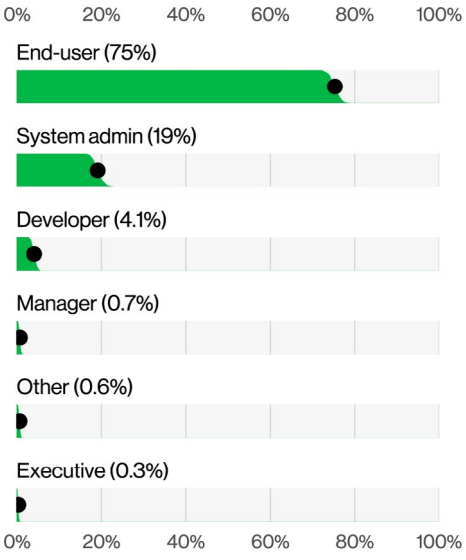

Figure 24. Top Internal actor varieties in breaches (n=1,000)

25. Yes, we know we discuss this in the "How to use this report" section, but we also know many of you don't read it.

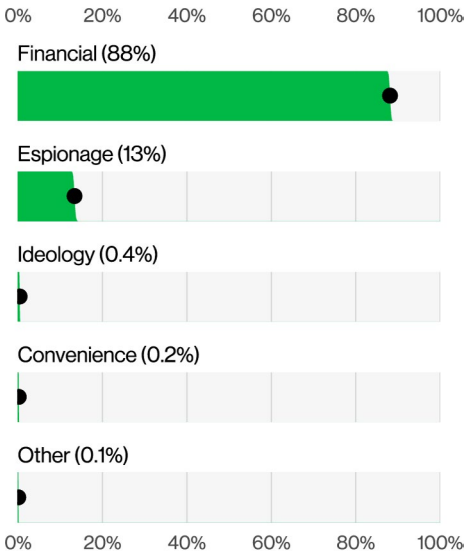

Figure 25. Top actor Motives in breaches (n=18,175)

As usual, these cases tend to be Espionage-driven, helping to bankroll foreign nation-states' more ambitious projects – whether that's advanced weapons programs or, who knows, maybe even chipping away at a few student loans.

Once we stop looking over the fence and turn our gaze inward, the internal actors driving risk are mostly familiar faces: End-users at 75% and System administrators at 19% (Figure 24). Internal actors may only account for 12% of breaches overall (both unintentional errors and deliberate actions), but when an employee acts maliciously, the blast radius can occasionally rival – or even exceed – that of an External attack. This is especially true when privileged access or sensitive systems are involved.

It's also worth calling out that many of these internal cases aren't cloak-and-dagger sabotage; they're often mistakes, Misconfigurations, or creative shortcuts and unapproved workarounds that simply backfire. These are all carried out without any malicious intent but with very real consequences.

But why, though?

Motives tend to mirror who is behind the attack. As Figure 25 shows, Financial gain remains the leading driver for cybercriminals (we'll pause while you recover from the shock).

The Espionage motivation is also present, albeit to a much smaller degree, as shown by the presence of state-affiliated actors in the data.

While most Espionage cases involve those groups, the same motive also surfaces among internal actors looking to walk off with proprietary data – whether to hand it to a competitor or to benefit themselves in some other manner. These kinds of cases are frequently the hardest to catch without specific offboarding processes in place for those employees who have access to the most sensitive data.

Actor categories[26]

External: External threats originate from sources outside of the organization and its network of partners. Examples include criminal groups, lone hackers, former employees and government entities. This category also includes God (as in "acts of"), "Mother Nature" and random chance. Typically, no trust or privilege is implied for external entities.

Internal: Internal threats are those originating from within the organization. This encompasses company full-time employees, independent contractors, interns and other staff. Insiders are trusted and privileged (some more than others).

Partner: Partners include any third party sharing a business relationship with the organization. This includes suppliers, vendors, hosting providers and outsourced IT support. Some level of trust and privilege is usually implied between business partners. Note that an attacker could use a partner as a vector, but that does not make the partner the Actor in this case. The partner has to initiate the incident to be considered the responsible party.

26. verisframework.org/actors.html

Yes, I'd love to help you write that malware.

The number of reports of malware and other hacking tools that leverage AI-assisted code[27] or elements of large language models (LLMs) in their workflows[28] has really grown in the last few months of 2025 and early months of 2026. Threat intelligence groups have been working hard to try to pin down this frontier-of-attack technology and to discover if there is something really novel going on.

Over the past couple of years, the DBIR has maintained a skeptic stance around the usage of those tools and whether they would move the needle in a way that could be measured by real incidents. Of course, experimentation happens, but how much would that actually change the likelihood an organization gets breached?

Would the use of AI by attackers even be notable or visible from an organization's perspective, especially considering the current scale and scope of cybercrime activity?

To try to answer some of those questions, we collaborated with Anthropic (developers of the Claude AI model) to understand their perspective on how threat actors are misusing their platform for nefarious cyberactivity. Anthropic's recent report[29] about the first documented case of a largely AI-executed state-sponsored espionage campaign was a watershed moment for understanding the potential and risks of the technology for many cybersecurity practitioners, your authors included. Building on this research, we wanted to understand the overall usage of LLMs by threat actors and its implications for defenders.

It's not depth of use – it's breadth of reach.

Anthropic's dataset covers 793 unique threat actors between Mar 2025 and Feb 2026. All 793 received enforcement action from the Anthropic Safeguards Team for violating the acceptable use policy and had sufficient behavioral data to analyze. Their queries spanned malicious cybersecurity topics, including malware development, capability building and tasking. Anthropic took this data and classified the behavior against the MITRE ATT&CK framework, which provides us a unique perspective to correlate this activity with what defenders can see from their perspectives. This included both cases of actors using the platform to write code to perform one or more of the techniques or leveraging it to complete the action in an agentic fashion.

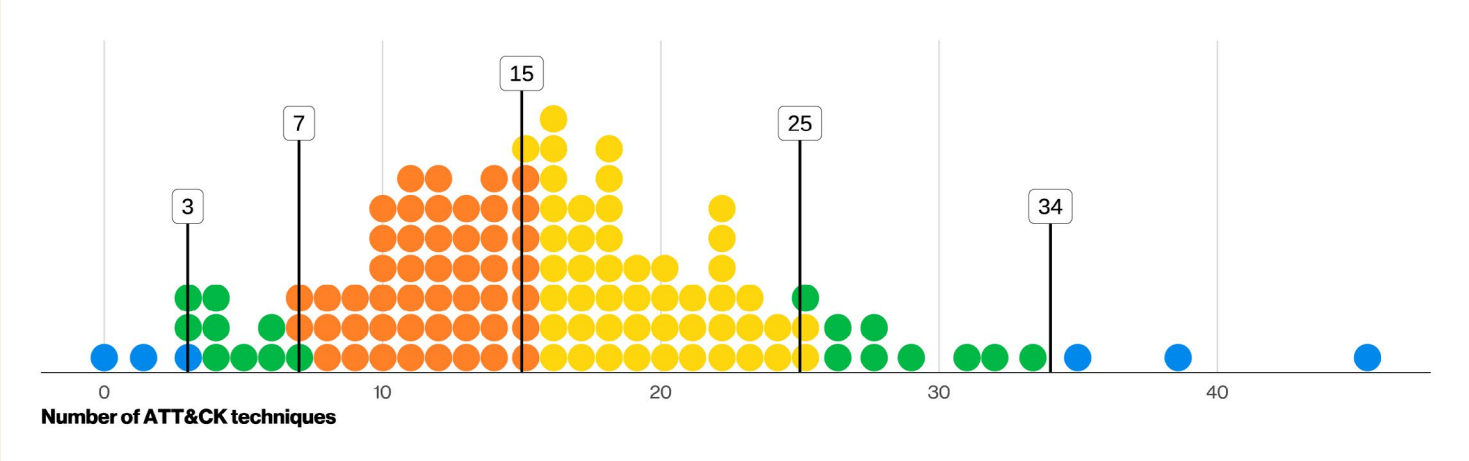

Figure 26. Distribution of unique ATT&CK techniques researched by actors on AI platform (n=793)

27. bitdefender.com/en-us/blog/businessinsights/apt36-nightmare-vibeware
28. cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use
29. anthropic.com/news/disrupting-AI-espionage

Figure 26 shows the distribution of techniques that were leveraged by each of those malicious actors. In the median case, actors sought AI assistance for around 15 distinct techniques across the MITRE ATT&CK. In the extreme cases, actors were querying for as many as 40 or 50 techniques, representing multisession campaigns where actors treated the platform as a co-developer across the full attack chain.

It's worth pointing out that these are extreme cases. Anthropic rates threat actor risk using signals such as number of detection hits, technical sophistication and which platform was used for the activity (e.g., code assistance plugin or the chatbot interface) – and less than 1% of those 793 actors fell into the High or Critical category, while 99% of actors fell into the Medium and Low Risk categories.

Here is how we break it down.

To start off, let's look at the reported techniques that fall under the Initial Access ATT&CK tactic and map them to our most common initial access vectors,[30] as they may explain some of the changes this past year or they might signal a future shift in pressure against organizations' defenses. Figure 27 has the details.

Although Exploitation of vulnerabilities and credential abuse match our previous analysis in the incident dataset, we can see Phishing in the lead aggregating 44% of the AI-assisted initial access vectors.

This lines up with some of our previous results from last year, where the increased use of AI-assisted text in malicious emails had doubled in relation to the previous years.

However, looking at our own incident dataset outside of Anthropic's data, the percentage of Phishing as an initial access vector has barely moved over the past few years, signaling that maybe AI assistance in this specific group of techniques is not increasing the success rate in the victim organizations that make up the DBIR's incident dataset. There could have been an increase of success rate in targeting individuals for fraud purposes that we are not able to measure, as they are out of scope of the analysis of this report. It can uplift less-experienced groups to a higher baseline level of proficiency in English, or the target language of their social attack, but that new baseline might not be enough for a higher success rate.

However, we note that 32% of initial access techniques are related to exploitation of vulnerabilities, a concerning finding given its growth trajectory in our incident initial access vector analysis. Leveraging AI coding assistance tools to create an exploit tool, change the language of the tool or discover new potential vulnerabilities is within reach with current AI coding assistance.[31]

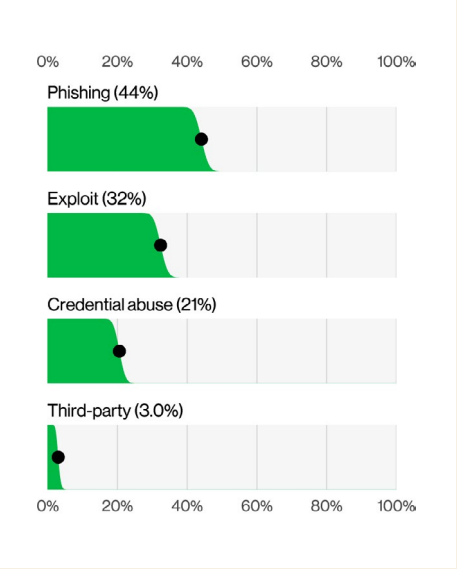

Figure 27. Generative AI-assisted techniques categorized as initial access methods (n=837)

30. As discussed earlier in the report in the "Results and analysis" section
31. We will discuss this recent development a bit more in a sidebar in the "VERIS Actions" section.

The shocking reason why that matters

A key question in understanding AI-enabled cyberthreats is whether attackers are using LLMs to execute well-documented techniques more efficiently, or to pursue techniques that are rarely seen in practice. If LLMs are lowering the barrier to techniques that are less documented and rare, defensive postures will need to catch up. To investigate this, the DBIR team developed a measure of technique rarity based on MITRE ATT&CK's catalog of known malicious software and offensive tooling.

MITRE publishes a list of known malicious software or offensive security tooling and the techniques that are leveraged by each of them.[32]

The hypothesis: The fewer existing tools that support a given AI-assisted technique, the rarer that technique would be in the wild.

This helps provide visibility into whether the attackers are using LLMs to implement techniques that are already well known and documented, such as Process Injection, or if they are leveraging LLMs for lesser-known techniques. It's worth noting again that since Anthropic's analysis is bound by MITRE ATT&CK, "rare" here refers to techniques that are operationally uncommon or difficult to execute within the existing taxonomy – not novel techniques altogether.[33]

Having said all that, Figure 28 shows the number of existing software examples per technique observed in abuse cases.

The median sits at 55, likely meaning most AI-assisted techniques already have dozens of known tools that implement them. In other words, the most common uses are well-trodden paths. The techniques with the most existing software examples show actors are likely outsourcing basic tasks to AI, such as file obfuscation or forensic cleanup, whereas the ones with fewer software examples demonstrate more creativity, with "Pre-OS Boot: The Unified Extensible Firmware Interface (UEFI)" and "Process Injection: Virtual Dynamic Shared Object (VDSO) Hijacking" as standouts. According to this classification, less than 2.5% of the techniques observed could be classified as rare (i.e., one known software example available or fewer).

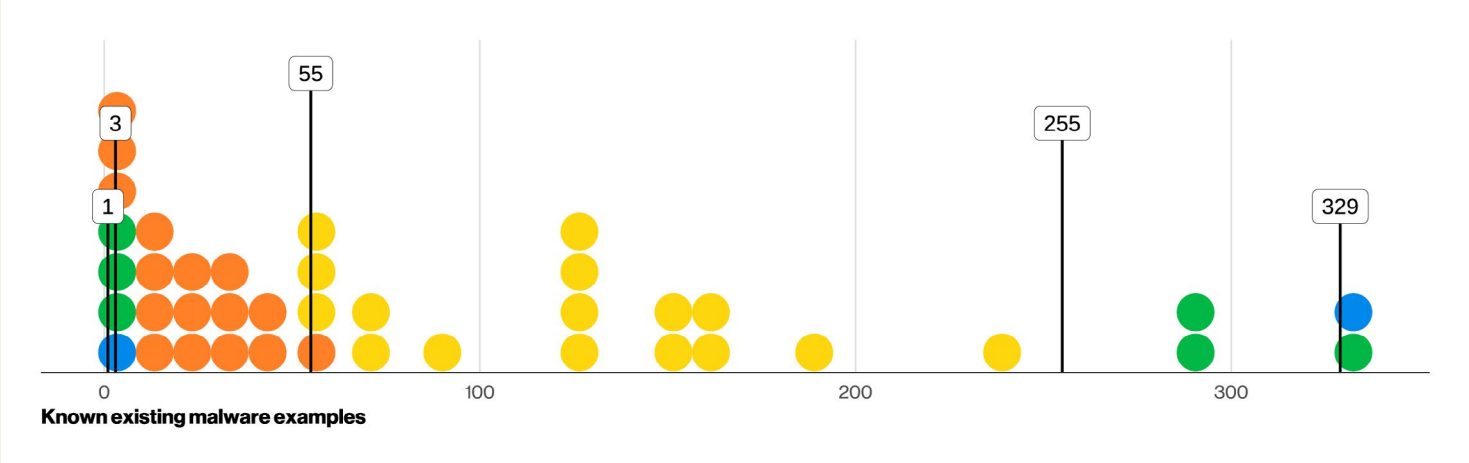

Figure 28. Distribution of known existing malware examples per ATT&CK technique observed (n=9,897)

32. attack.mitre.org/software
33. The world is not ready for a groundbreaking, quantum-enabled, blockchain-powered metaverse attack.

The takeaway from our dataset is that AI's primary impact is currently operational: automating and scaling techniques defenders already know how to detect, not yet unlocking these novel or rare attack surfaces – which means defensive postures don't need to be reinvented today, but they do need to keep pace with faster, more adaptive execution. But who knows? Given the rate of change in AI capabilities, this assessment might be obsolete by the time this report is finally published.[34]

Can I help you hack another company?

The current state of our dataset reinforces what we're seeing across the threat landscape: AI is primarily accelerating the operationalization of well-known, documented techniques – lowering the barrier to execute what was once out of reach for less-sophisticated actors. Anthropic's previous reporting from Nov 2025 actually demonstrates that the real sophistication of actors has historically come from the behavior they do around the MITRE technique assistance they get from AI. The more novel cases include combining or chaining together multiple stages of the attack or taking more agentic approaches to the attack, where the agent makes executive decisions about the targets.

Still, this raises the tide for all attackers, and the baseline of what can be achieved with relative low cost is broader. In a world where the security poverty line[35] exists, any industrialization, even of simple techniques, could make the gap between the cybersecurity haves and have-nots even wider.

Some of the top variables we need to keep track of here as an industry are how many of the actors start to present a higher-risk profile, how many are getting assistance in a greater variety of techniques, and which ones point to more novel behavior in the software ecosystem where our ability to mitigate and detect are not as well developed.

34. Who is reading this report on their transparent, foldable phone in 2030?
35. Shout out to Wendy Nather and her seminal paper, "T1R Insight: Living Below the Security Poverty Line": web.archive.org/web/20140203193523/https:/451research.com/t1r-insight-living-below-the-security-poverty-line

VERIS Actions

In the "VERIS Actors" section, we covered the "who" behind the attacks in our dataset. Here in the "VERIS Actions" section, we shift to the "how" – the methods those threat actors used to ruin your day. You may think of actions in cyberattacks like fashion trends: They come in, they go out and just when you think you've seen the last of them, they resurface on your runway.[36] In any case, we wouldn't mind if a few of these techniques went the way of velour tracksuits and never came back.

As you will see in other parts of the report, particularly in the "Industries" section, the top-level action types of Hacking and Malware often appear in relatively equal measure as illustrated in Figure 29. Hacking actions showed up in 64% of breaches this year, and a corresponding Malware action appeared in 63% of breaches.

This makes sense because all that malware doesn't just magically appear on a system – it usually needs a little help to get its foot in the door, and this assistance frequently comes from some kind of Hacking action.

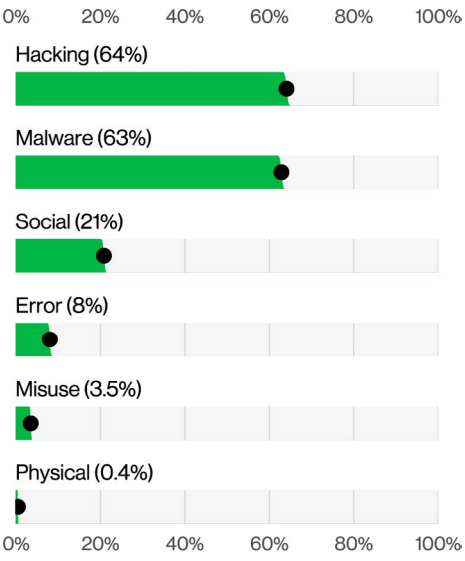

Figure 29. Actions in breaches (n=22,624)

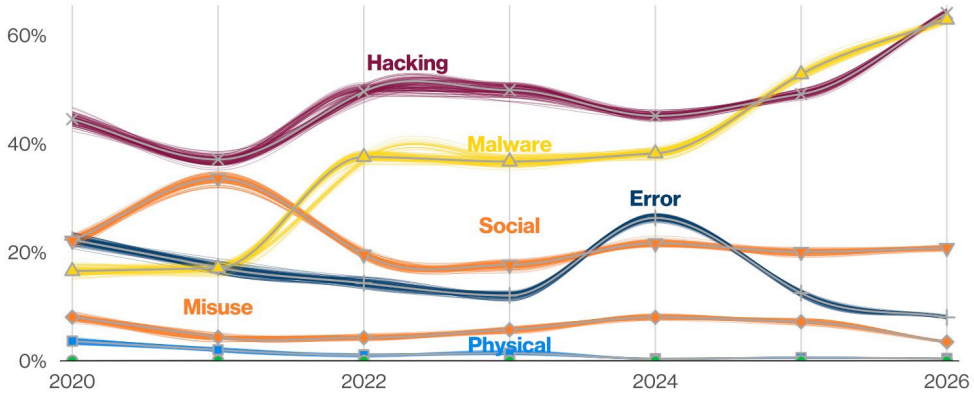

Figure 30. Action varieties over time (n for 2026 dataset=22,624)

Once inside, threat actors often pair Malware with additional Hacking actions to wander laterally around the victim's network as if they own the place. As Figure 30 shows, both Hacking and Malware have been rising in lockstep over the past few years. Much of that consistent rise has been driven by the ever-present Ransomware attacks we see, well, more or less everywhere. Meanwhile, Social attacks played a role in a comparatively low 21% of breaches.

The top action varieties in breaches (Figure 31) confirm that Ransomware is still the yoga pants of cybersecurity – ubiquitous, stubbornly popular and appearing in unexpected places near you. Nearly half (48%) of breaches in this report feature Ransomware somewhere in the attack chain. The sizable Other category in second place isn't hiding anything exotic; it's just a parking spot for Hacking and Malware actions we know occurred but couldn't cleanly classify due to limited visibility.

The Use of stolen credentials holds steady at 36%, continuing its long-running role as an attacker favorite. Exploitation of vulnerabilities, however, is having a breakout season. After hitting an all-time high of 18% in last year's report, it has doubled to 32% of breaches.

This highlights how unpatched or otherwise exposed systems are still rolling out the red carpet to cybercriminals rather than being the hardened entry points they were thought to be.

Backdoor or command and control (C2) functionality appears in 16% of breaches, only a modest increase from the previous year's 14%.

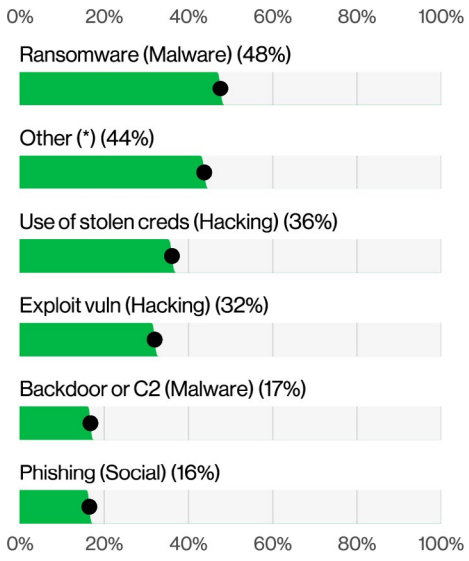

Figure 31. Top action varieties in breaches (n=19,550)

36. Admit it, you missed parachute pants so much, you're secretly delighted we brought them back. You're welcome!

One reason this number isn't higher comes into focus in Figure 33: Threat actors seem to be dialing back their reliance on traditional tools such as Cobalt Strike as defenders get better at detecting and blocking them. Instead, attackers are increasingly blending in with normal operations by using the same access paths as legitimate users – such as desktop-sharing tools and virtual private networks (VPNs) – which helps explain the noticeable rise of those vectors in recent incidents.

Finally, a quick word on the top action varieties in incidents: They largely track what we see in the breach data, with a notable bit of extra flair in the form of DoS attacks (Figure 32). In other words, various forms of Hacking, Malware and Social tactics tend to figure prominently here, as well. However, incidents add another layer of woe with disruption-focused activity rather than simply data compromise. DoS attacks are less about stealing information and more about kicking the chair out from under Availability, reminding organizations that not every incident is about data loss.

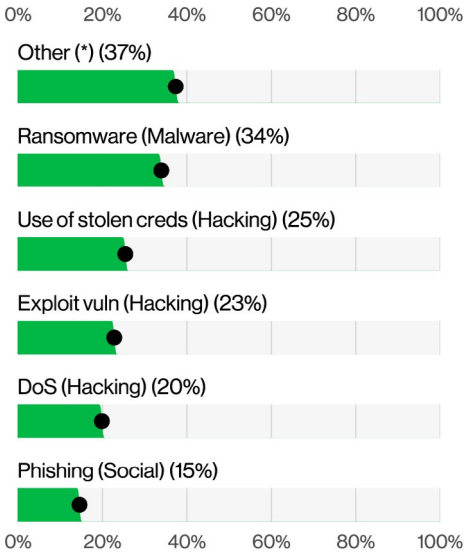

Figure 32. Top action varieties in incidents (n=27,765)

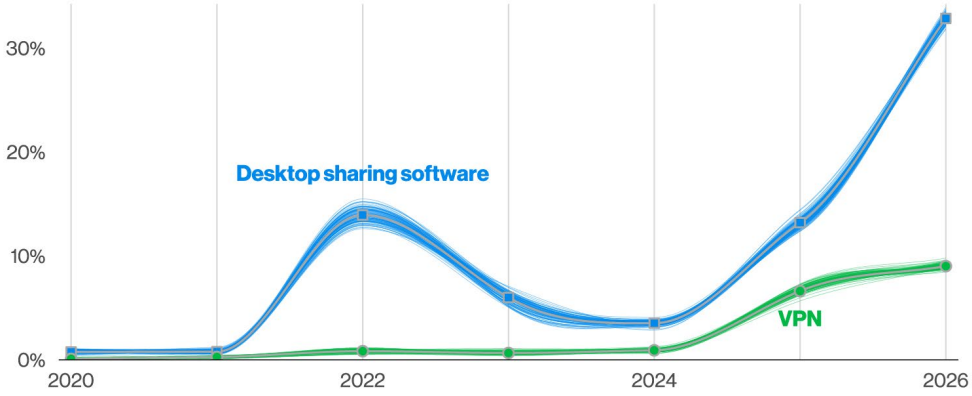

Figure 33. Select action varieties over time (n for 2026 dataset=18,175)

Action categories[37]

Hacking: attempts to intentionally access or harm information assets without (or exceeding) authorization by circumventing or thwarting logical security mechanisms.

Malware: any malicious software, script or code run on a device that alters its state or function without the owner's informed consent.

Error: anything done (or left undone) incorrectly or inadvertently.

Social: employ deception, manipulation, intimidation, etc., to exploit the human element, or users, of information assets.

Misuse: use of entrusted organizational resources or privileges for any purpose or manner contrary to that which was intended.

Physical: deliberate threats that involve proximity, possession or force.

Environmental: not only includes natural events such as earthquakes and floods but also hazards associated with the immediate environment or infrastructure in which assets are located.

37. verisframework.org/actions.html

Exploring the weakness within

This report has already spent quite a lot of pages talking about the prevalence of vulnerabilities and the struggle to patch them, but there is another interesting analysis angle that can shed some light on the difficulties we are facing.

This subsection was inspired by the great work from the MITRE CWE[38] team, where they rank the CISA KEV CVE vulnerabilities discovered each year by the weaknesses they are based on.[39] Every single vulnerability has an underlying technical raison-d'être, and those are captured in the weakness enumerations of the CWE. For example, something like a Remote Code Execution (RCE) vulnerability could be caused by CWE-122, a "Heap-based Buffer Overflow" or CWE-416, a "Use After Free."

The CWE focuses on the root cause of the vulnerabilities, so this is a sensible and informative way of grouping together vulnerabilities that otherwise would make little sense to aggregate. They focus on different software, and the vulnerabilities were discovered at different times by different security researchers—you get the picture.

Instead of looking at which vulnerabilities were discovered in 2025, let's have a look at the ones that were the most detected in our vulnerability management dataset throughout the analysis period of the DBIR and find the weaknesses most commonly found therein.

Table 1 only shows the top five, but all of those weaknesses have been among vulnerabilities detected in more than 75% of the organizations in our vulnerability management dataset.

CWE	Percentage
CWE-125 - Out-of-bounds Read	79%
CWE-122 - Heap-based Buffer Overflow	77%
CWE-416 - Use After Free	77%
CWE-73 - External Control of File Name or Path	77%
CWE-843 - Access of Resource Using Incompatible Type ('Type Confusion')	76%

Table 1. CWE frequency of detected CISA KEV CVEs per organization (n=12,208)

CWE category	Percentage
CWE-1399 - Memory Safety	89%
CWE-1396 - Access Control	85%
CWE-1416 - Resource Lifecycle Management	80%
CWE-1407 - Improper Neutralization	77%
CWE-1404 - File Handling	77%

Table 2. CWE category frequency of detected CISA KEV CVEs per organization (n=12,208)

In fact, to find a weakness present in less than 70% of organizations, you would need to go beyond the top 10. This is already very informative, and almost all of the top five seem related to memory safety vulnerabilities.[40]

We can summarize this further by rolling up the CWEs to their top-level categories.[41] The results in Table 2 confirm our fears. A staggering 89% of organizations had to patch vulnerabilities associated with Memory Safety. Memory Safety! "Smashing the stack for fun and profit"[42] is from 1996, 30 years ago! What are we still doing here? If this vulnerability were a person living in the U.S., they would not only be able to drive but could also rent a car without paying a "young rental fee."

This is why so much of the focus of Secure by Design[43] initiatives has been centered around the elimination of classes of bugs like those and, specifically in this case, the advocacy of using memory-safe languages for developing software.

As a bonus, we thought it would be interesting to take a look at the other side of this equation, courtesy of one of our new contributors that provided us with a very robust dataset of the detection of code flaws during the software development life cycle. One of the metrics we can extract is the time of resolutions by categories of flaws.

By calculating the survival analysis for every single CWE grouping of flaws found, we can get a measurement in months of how long it takes to correct the code and re-submit for testing of 50% of the flaws associated with each CWE.

38. Not a typo, we do mean the Common Weakness Enumeration this time.
39. cwe.mitre.org/top25/archive/2025/2025_kev_list.html
40. We can only speculate as to whether or not Rust would have prevented these from happening.
41. More about those categories, from the MITRE CWE itself: cwe.mitre.org/data/definitions/1400.html
42. A seminal paper published in Phrack 49 that introduced the concepts of buffer overflow, a common memory safety exploit technique: phrack.org/issues/49/14
43. cisa.gov/resources-tools/resources/secure-by-design

By arranging those findings in a distribution for each of the CWE categories we used earlier, we can get a good sense of how hard it is to resolve our top three CWE categories from the previous figure in practice for companies with mature Software Developer Life Cycle (SDLC) loops.

As Figure 34 shows, our top three CWE categories have a median 50% survival rate of between six and seven months.[44] Re-testing cycles can potentially vary a lot depending on the application or organization's SDLC processes, but this result highlights how costly (or plain time consuming) resolving the code flaws can be if fixing them is taken seriously by the developer. If you are curious, the worst median 50% survival time for a CWE category is for "Improper Input Validation" – another nigh-unforgivable[45] weakness category – sitting at just a bit over 13 months.

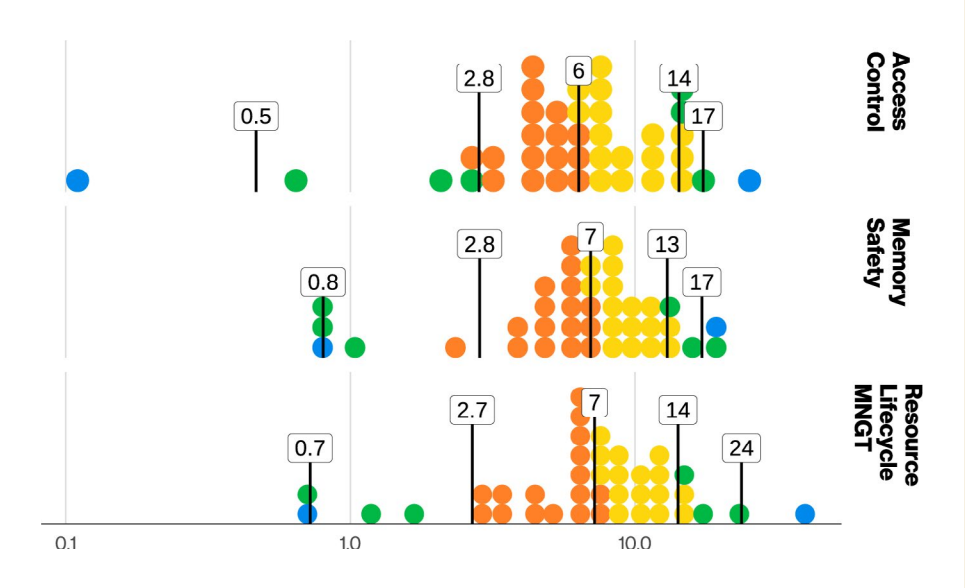

Figure 34. Distribution of time to remediate 50% of codebase weaknesses (in months)

Artificial strength training[46]

It is worth highlighting the recently declared success in the usage of GenAI platforms to discover large numbers of new vulnerabilities in code bases, and we wonder how this will actually change the calculus of discovery and resolution of code flaws ahead of shipping software. Static Application Security Testing (SAST) and Dynamic Application Security Testing (DAST) have distinct capabilities to find different types of weaknesses, and maybe the GenAI solutions, when properly tuned and productized, would eventually lean toward SAST due to the whole "analyzing the source code in bulk" angle. It's hard to conjecture, though, without objective data at hand.

As solutions pairing those detection capabilities with suggestions on how to fix them emerge, will those remediation capabilities be equally effective across all vulnerability types, and would their detection strengths align with the areas where those automated fixes are the most needed. We look forward to evaluating this as more data becomes available, but the expectation is that a model from an AI frontier lab with focused post-training and fine-tuning would likely[47] achieve good results. Still, "Quis custodiet ipsos custodes?"[48] comes to mind.

There is a lot of potential for changing the landscape of the "code flaw to CVSS 10.0 vulnerability pipeline" if the right types of weaknesses are targeted by those solutions. In the meantime, we'll continue to promote the foundation and importance of Secure by Design (with AI assistance when available) and tried and true processes for software testing and remediation, even if we are still having issues with 30-year-old classes of vulnerabilities. Outside your own development pipeline, one recommendation is to prepare for a large number of patches from coordinated disclosures of AI-augmented vulnerability discovery, while making sure to inventory and minimize your internet-facing footprint.[49] The zero-days will keep being discovered and weaponized, regardless if fueled by tokens, by coffee or most likely a mix of both.

The DBIR team wishes that there was more available and verifiable data on what the future defensive landscape leveraging those technologies will look like. However, as reductive as "only a good guy with AI can stop a bad guy with AI" may sound, practitioners should spend time exploring this AI frontier to help improve their defenses when possible. The speed of those advancements really puts our only-publishes-once-a-year DBIR framework to the test.

44. We know how it looks, but we did not mean to make a six-seven joke here.
45. Speaking of influential papers, it's never a bad time to reference "Unforgivable Vulnerabilities" by Steve Christey Coley from 2007: cwe.mitre.org/documents/unforgivable_vulns/unforgivable.pdf.
46. This section was written in Feb 2026, roughly a couple months before the advancements in frontier models by AI companies. It has been updated to provide actionable recommendations on a topic that becomes less theoretical by the minute.
47. And at least on the flaw detection and exploit development front in Apr 2026, the results have the foremost experts in vulnerability discovery and disclosure on the edges of their seats.
48. "Who watches the watchmen?," a rhetorical question way older than the graphic novel itself
49. For even more in-depth recommendations and discussions of impact, we suggest taking a look at a recent briefing put together by the Cloud Security Alliance: labs.cloudsecurityalliance.org/mythos-ciso.

VERIS Assets

Assets are ultimately the targets of the attacks. This is where the proverbial rubber meets the road in terms of whether the ploy of the adversaries is successful in furthering their goals. They are also the places where your defenses can make the most difference in stopping an attack in its tracks. Understanding which assets are targeted is critical to drive the types (and placement) of controls an organization needs to deploy.

As in prior years, we break the data down in multiple ways, including which assets are targeted most per industry (see the introduction in the "Industries" section). This provides defenders with a basic road map to link the attacks targeting your industry to the assets they are likely to hit. It works well as a tabletop exercise to determine if your organization's controls as they currently exist would be able to detect and respond to the top attacks in your industry.

Figure 35 shows the ranking of the VERIS Assets in breaches. The Server and Person assets are in consistent rankings with last year, but the Network asset has seen quite an increase this year.

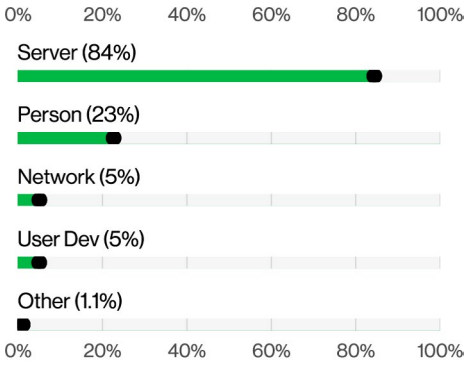

Figure 35. Asset types in breaches (n=20,862)

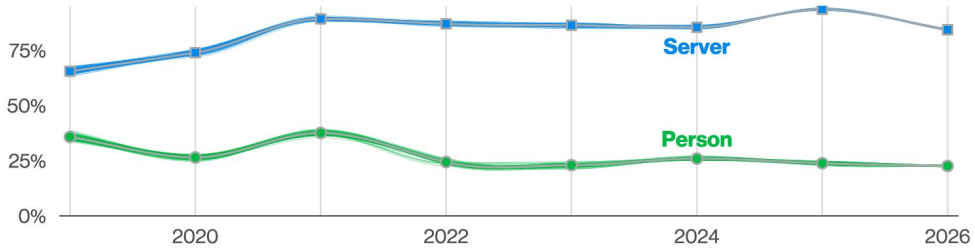

Figure 36. Select top assets in breaches over time (n for 2026 dataset=20,862)

It accounted for just 1.5% of breaches in last year's report, but it is now tied with User devices, accounting for 5% of breaches where the asset was known. If you look at this chart on the 2025 report, you will find a much smaller value for Network (0.3%). However, we changed the way we code cases involving remote access from a Server to Network to better align with those devices' edge of network role. It is still a significant rise in percentage for this variety of asset.

Last year, User devices were targeted in 8% of attacks, putting them ahead of Network assets. This year, they have fallen to slightly below Network assets, but the difference between them is statistically negligible. Suffice it to say that Network and User device assets are similarly targeted in this year's dataset.

The pairing of Server and Person has dominated the asset landscape since at least 2019, with nothing else coming close to those two asset classes, as Figure 36 shows. However, there is still a significant difference in frequency between the two. This highlights how prevalent Social actions, which compromise the Person asset, remain despite ongoing efforts to reduce this risk through training and awareness.

Does this mean that your people will always fall for social attacks? No, but it does mean that you should make it as painless as possible for them to notify you when they do, so you have a chance to contain the damage before it escalates.

Figure 37 shows which Asset varieties we see most frequently targeted in these attacks. The top three are the same as last year. Unsurprisingly, web servers and mail servers are prime entry points for criminals attempting to break into an organization's infrastructure. The Person asset is also a tried-and-true target for Social attacks, as mentioned above.

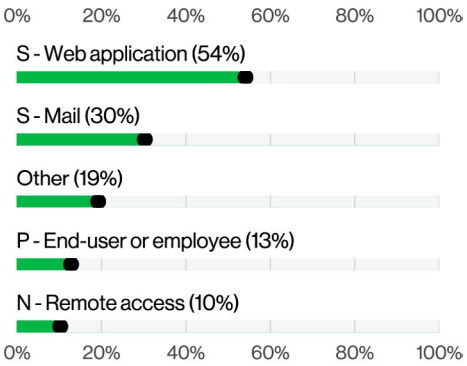

Figure 37. Asset varieties in breaches (n=9,723)

OT or not OT? That's the question.

Information about incidents in operational technology (OT) equipment is often sparse and has not grown enough in the past few years for us to do a deep dive of any sort.

However, there are almost 800 OT-adjacent incidents we have recorded this past year on NAICS 21 (Mining, Quarrying, and Oil and Gas Extraction) and NAICS 22 (Utilities), which may give some insights on how these critical sectors are being impacted in the threat landscape.

On the topic of OT-adjacent networks, there is one specific threat we would like to point out based on mass internet scanning and network telemetry around end of life (EOL) internet-facing cellular routers in NAICS 21 and 22 companies. Those EOL devices are either vulnerable to configuration oversights – such as default passwords in internet exposed management interfaces – or unpatched publicly known vulnerabilities.

Those devices were often put in place to enable Internet of Things (IoT) sensors and connectivity to remote outposts and are being repurposed as operational relay boxes (ORBs) to be used as proxies for threat actor attack campaigns.[50] For example, our internet scan analysis validated the existence of between 45,000 and 50,000 EOL wireless modem devices with a publicly accessible ACEmanager interface between the months of June and October of 2025.

These are not easy-to-replace devices even for the most well-intentioned organizations, and that makes network segmentation one of the simplest solutions to help mitigate the issue ahead of a potential replacement of the EOL modems. Recent technologies such as 5G network slicing[51] and private mobile networks would be something worth taking a look at if this is a scenario that concerns your organization.

Asset categories[52]

Server: a device that performs functions of some sort supporting the organization, commonly without end-user interaction. This is where all the web applications, mail services, file servers and all that magical layer of information is generated. If someone has ever told you "the system is down," rest assured that some Servers had their Availability impacted. Servers are common targets in almost all of the attack patterns, but especially in our System Intrusion, Basic Web Application Attacks, Miscellaneous Errors and Denial of Service patterns.

Person: the folks (hopefully) doing the work at the organization. No AI chatbots allowed. Different types of Persons will be members of different departments and will have associated permissions and access in the organizations stemming from their roles. At the very least, they will have access to their very own User device and their own hopes and dreams for the future. Person is a common target in the Social Engineering pattern.

User device: the devices used by Persons to perform their work duties in the organization. Usually manifested in the form of laptops, desktops, mobile phones and tablets. These are common targets in the System Intrusion pattern but also in the Lost and Stolen Assets pattern. People do like to take their little computers everywhere.

Network: not the concept, but the actual network computing devices that make the bits go around the world, such as routers, telephone and broadband equipment, and some of the traditional in-line network security devices, such as firewalls and intrusion detection systems. Hey, Verizon is also a telecommunications company, OK?

Media: precious distilled data in its most pure and crystalline form. Just kidding, mostly thumb drives and actual printed documents. You will see the odd full disk drive and actual physical payment cards from time to time, but those are not common.

50. This repurposing of assets to attack other victims is referred to in VERIS as a Secondary actor motive.
51. What if your cell tower acted like a network switch full of virtual local area networks (VLANs), or more accurately virtual wide area networks (VWANs)?
52. verisframework.org/assets.html

VERIS Attributes

We've covered the who, how and what of data breaches; now it's time to examine the "so what" of the incidents. The Attributes capture what impact an incident actually had on the Asset and are rooted in the well-known CIA triad: Confidentiality, Integrity and Availability. In plain terms: Did the attack expose data, alter the asset in some way, take the asset offline or encrypt it – or some combination of all three?

Because we define a breach as any incident where Confidentiality is compromised, every case labeled as a "breach" in this report includes at least some form of data disclosure. But depending on what happened, Integrity and Availability can easily get pulled into the mess, as well. A single event can leak data, corrupt records and knock systems over, all in one go. For more information, please see the definition of each in the callout.

Attribute categories[53]

Confidentiality: refers to limited observation and disclosure of an asset (or data). A loss of confidentiality implies that data was actually observed or disclosed to an unauthorized actor rather than endangered, at-risk or potentially exposed (the latter fall under the attribute of Possession or Control[54]). Short definition: limited access, observation and disclosure.

Integrity: refers to an asset (or data) being complete and unchanged from the original or authorized state, content and function. Losses to integrity include unauthorized insertion, modification and manipulation. Short definition: complete and unchanged from original.

Availability: refers to an asset (or data) being present, accessible and ready for use when needed. Losses to availability include destruction, deletion, movement, performance impact (delay or acceleration) and interruption. Short definition: accessible and ready for use when needed.

Figure 38 looks at how Attributes were affected across all incidents in this year's dataset. 82% of incidents showed confirmed data disclosure – meaning someone who shouldn't have access either viewed or downloaded confidential information. Integrity took a hit in 64% of incidents, ranging from phishing campaigns that manipulate users into clicking malicious links to the unauthorized installation of software on victim assets. Availability was impacted in 53% of incidents, driven by events such as successful DoS attacks or ransomware that led to the encryption of the victim's data, leaving systems offline or their contents unusable.

The share of incidents in which the Interruption variety of availability was affected rose significantly in this year's report. Two main factors are at play here. First, VERIS is not static – the schema evolves as attack types change and as we refine how we capture incident data. Historically, ransomware was primarily coded under the Obscuration variety of availability, since the data was encrypted. We did not consistently capture the Interruption aspect, or in other words, the system downtime that often accompanies a successful ransomware attack. This year's coding more fully reflects that operational disruption. Second, we are seeing (and coding) more large-scale ransomware outages, such as the Apr 2025 incident affecting Marks & Spencer.[55] In that case, online sales, stock tracking and reordering, and even the electronic monitoring of refrigeration was disrupted for weeks, resulting in an estimated £300 million in losses from prolonged outages.

53. verisframework.org/attributes.html
54. en.wikipedia.org/wiki/Parkerian_Hexad
55. blackfog.com/marks-and-spencer-ransomware-attack

Naturally, these high-impact events exert quite an influence on the dataset and help push the overall percentage of incidents involving the Interruption variety of availability upward.

Figure 39 lays out the top data varieties in breaches, and Internal data is the most commonly stolen at 67%. That's not exactly a head-scratcher when you remember that "Internal" mostly means emails, plans and reports – the kind of material you'd expect to be lying around once an attacker strolls in via stolen credentials or an unpatched vulnerability. It would almost be more surprising if they didn't take a peek.[56]

Credentials appeared in 28% of breaches and are relatively self-explanatory – and entirely on brand – given the password habits of many organizations, as we've been lamenting for 19 years.

Personal data (names, addresses and phone numbers) was taken in 23% of breaches, with sensitive personal information, such as Social Security numbers getting compromised in 1.7% of breaches.

In short, the data shows that once attackers are inside, they can (and do) go in many directions in an organization's environment, and whether they are looking for something specific or just looting as they go, much of the data in their path may be at risk. Between the persistent headache of stolen credentials and ever-rising cost of operational downtime, the "so what" of an incident seems to be getting more expensive by the year. Perhaps the takeaway is that you should not only be concerned about what they will see but also what they may break.

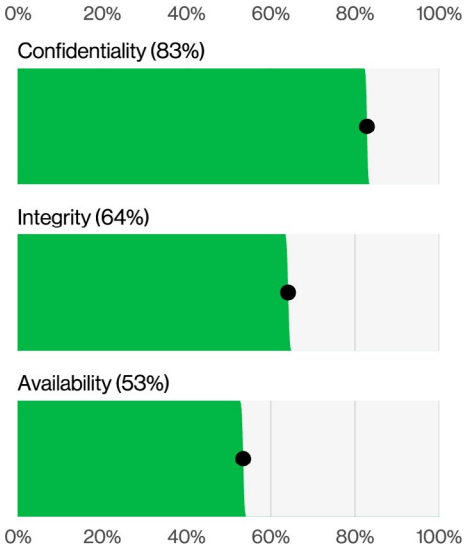

Figure 38. Attributes in incidents (n=31,850)

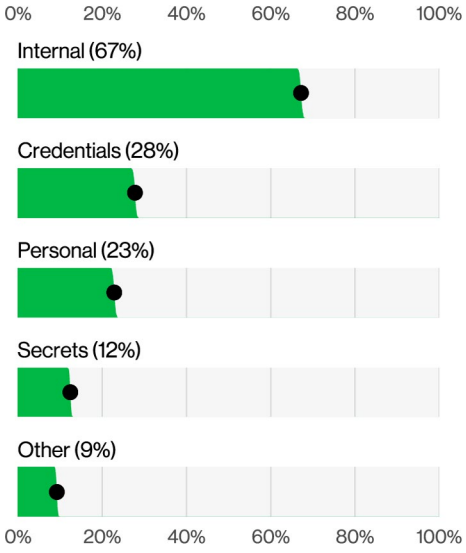

Figure 39. Top Data varieties compromised in breaches (n=19,538)

56. Sadly, the adage does not say "Curiosity killed the criminal."

Incident Classification Patterns

/03

Introduction

After watching case after case cross our desks, it's hard to ignore that most security incidents are essentially reruns of the same plot, with only the occasional ad-lib by the attackers. The same traits often kept appearing together, and those recurring combinations gave us just enough structure to tame the chaos into a sensible set of categories. Since most people find it easier to reason about complex ideas when they can be placed into clearly labeled containers, those Incident Classification Patterns became our organizing framework – and they have remained with us ever since, bolstered along the way by updated machine learning models that handle much of the classification work.

These patterns are still the primary lens we use to group incidents that share similar characteristics into categories that are easier to understand, explain and recall. They are grounded in the 4As of VERIS – Action, Actor, Asset and Attribute – which together capture how something happened, who was involved, what was affected and in what way.

Using that structure, we arrive at seven Incident Classification Patterns that you will see referenced throughout this section. Each of the pattern sections will include the CIS Security Controls[57] relevant to them.

Returning readers may also notice some editorial decisions about what earns a full discussion versus what is summarized at a glance. Certain patterns are remarkably consistent from year to year, both in terms of frequency and in the defenses that matter most. In those areas, the story does not change much: The same tried-and-true practices often continue to have the greatest impact, and the data does not warrant pages of new commentary. Where the data does surprise us through sharp increases, notable declines or meaningful shifts in how attacks play out, you will see those changes called out and examined in more detail in the pages that follow.

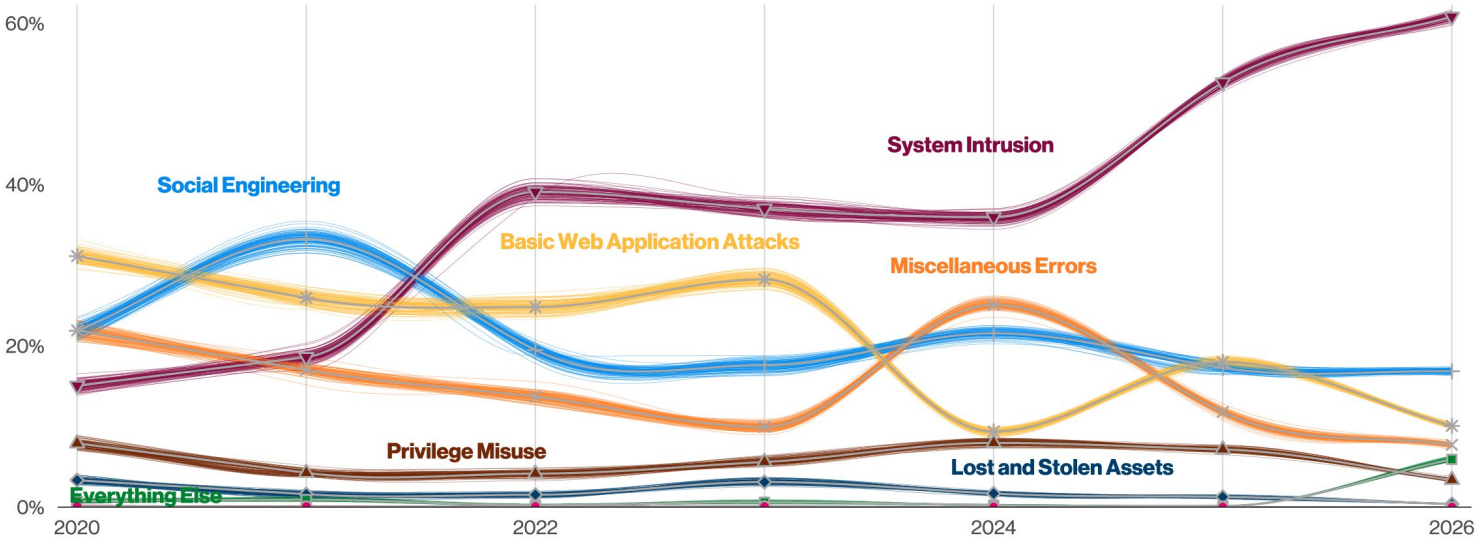

Figure 40. Patterns over time in breaches (n for 2026 dataset=22,624)

57. cisecurity.org/controls

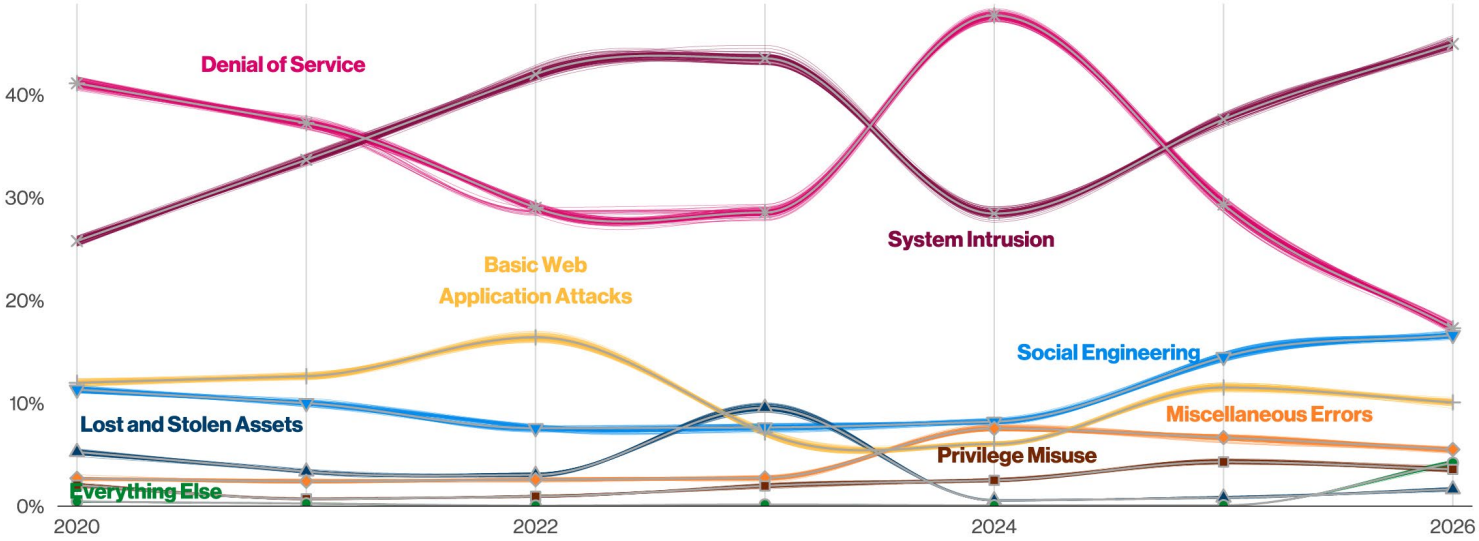

Figure 41. Patterns over time in incidents (n for 2026 dataset=31,860)

Lost and Stolen Assets

Frequency	525 incidents, 88 with confirmed data disclosure
Threat actors	Internal (88%), External (13%), Multiple (1%) (breaches)
Actor motives	Financial (79%–100%) (breaches)
Data compromised	Personal (96%), Internal (43%), Other (21%), Sensitive Personal (18%) (breaches)

Summary

This year, we saw a sharp uptick in the number of incidents involving Lost and Stolen Assets. However, 82% of those were considered breaches in the 2025 report, and in this year's report, that number is only 17%. The fact that the number of confirmed breaches this year is so much lower than the number of incidents may be attributable in part to the victim organization being unable to confirm if the data on the asset was accessed.

What is the same?

Personal data remains the primary data type at risk in these cases. This pattern remains largely one of insiders making mistakes and losing track of their devices, that being roughly four times more likely than the assets in question "growing legs" and wandering off with the help of financially motivated thieves.

System Intrusion

Summary

Threat actors are leveraging trusted applications, such as Remote Monitoring and Management (RMM) software, stolen credentials and exploits to monetize their access via Ransomware.

What is the same?

Ransomware continues to be the driving force behind the growth of System Intrusion.

Frequency	14,309 incidents, 13,758 with confirmed data disclosure
Threat actors	External (100%) (breaches)
Actor motives	Financial (88%), Espionage (12%) (breaches)
Data compromised	Internal (93%), Credentials (26%), Other (20%), Secrets (13%) (breaches)

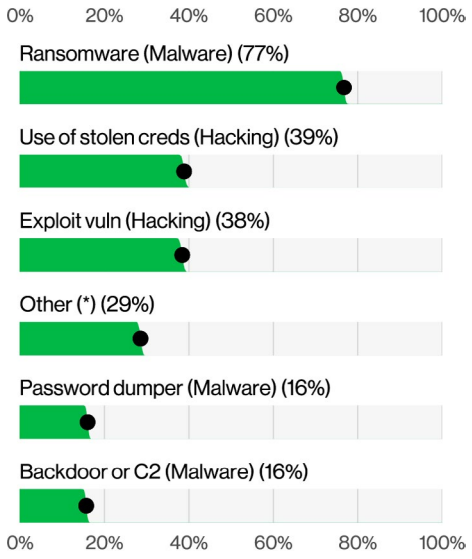

Figure 42. Top Action varieties in System Intrusion breaches (n=12,006)

While a good number of our breach patterns include some degree of intrusion being inflicted on systems of one sort or another, this pattern gets its name from the more complex and involved breaches. In these cases, determined external actors are leveraging a combination of Malware (and other software) and their knowledge of Hacking techniques to compromise our best-guarded data. This pattern has been our top breach pattern since 2022 and currently accounts for 60% of our breaches. It is difficult to overstate the contribution that the growth of Ransomware-like attacks has made to this pattern over the last few years.[58]

Intrusive techniques

Figure 42 describes the top Action varieties for the System Intrusion pattern. They provide a demonstrable cross-section of what types of attacks defenders are up against.

Unsurprisingly, Ransomware is at the top of the list and is seen in 77% of the breaches in the System Intrusion pattern. As we have pointed out in previous reports, these attacks represent the main method of monetization used by criminals in this pattern. Conversely, we see more of an equal split in the common initial access vectors Use of stolen creds and Exploit vuln (both at 39%).[59]

To a lesser extent, we see the more general lateral movement techniques, such as Scan network and the ever-popular Password dumper, being used in order to find and extract credentials from compromised systems. Look forward to a longer discussion of privilege escalation techniques – and their mitigations – later in this section.

Even though the actions themselves are in line with previous reports, the vectors used to achieve them show some interesting variation from last year. In Figure 43, we see a certain degree of evolution in tactics, with a growth in targeting Web applications and Desktop sharing software. This can be partially attributed to more exploited vulnerabilities against these types of services, while VPN maintains a similar level as last year, around 16%. It's worth noting that those vectors are also prime targets for credential abuse-related actions, and that contributes to their lasting presence over the years.

58. We wonder how future cybersecurity history will refer to this era. We suggest the Ransoming Twenties!
59. It's worth checking out our discussion of initial access vectors this year in the "Results and analysis" section.

Unapproved admin(ware)

One of the key things attackers enjoy doing is controlling a system remotely, which coincidentally is also something most administrators are also fond of.[60] Unfortunately, this coincidence works against our established technology and security infrastructure strategies.

In many cases, attackers leverage the typical hacker-style toolkit while also bringing in a C2 framework to manage their remote persistence agents on victim systems and to carry out activities such as dumping passwords and scanning local networks.[61]

However, Figure 44 presents a clear increase in threat actors leveraging legitimate RMM software to orchestrate their attacks. This year shows a sharp increase in the prevalence of those techniques, with a relative growth of 240% over the previous year. When coupled with the 27% decrease of the Backdoor or C2 action, it solidifies the belief that defenders have more than a few war stories to share relating to these types of scenarios.[62]

Those tools, while not chock full of hacking scripts, do provide actors with a remote session that they can easily leverage and deploy to an environment. They often have the added bonus of not requiring the actor to set up additional infrastructure. Instead, they are conveniently included in the application and network whitelists of the organization deploying them. At the end of the day, what is a threat actor but an unapproved (and malicious) administrator?[63]

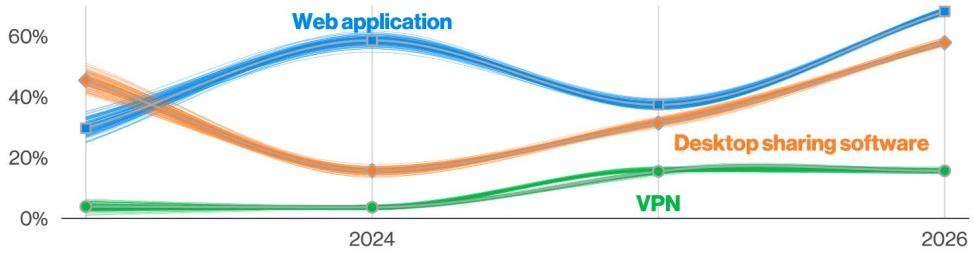

Figure 43. Select Action vectors in System Intrusion breaches (n for 2026 dataset=6,678)

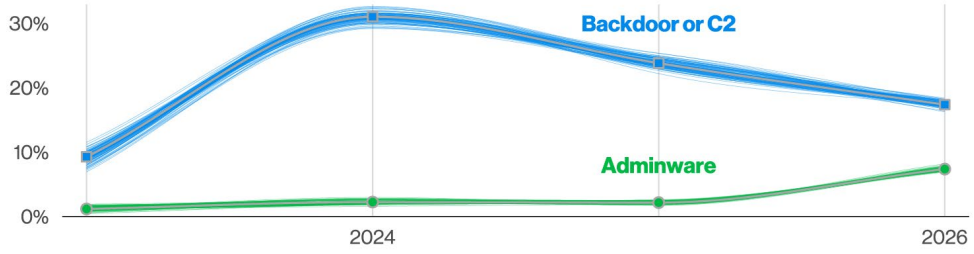

Figure 44. Select Malware varieties over time in System Intrusion breaches (n for 2026 dataset=10,828)

The invisible hand of the ransom market

Those who read the DBIR regularly are no doubt weary of being told how Ransomware has continued to grow over the past few years. This year's growth has not been as dramatic as last year (for which we are grateful), but in 2025, 48% of all the breaches analyzed had a Ransomware action involved, as discussed previously in the "VERIS Actions" section of this report. Even with this continual growth, there's an opposing trend we have been reporting around the decrease in the success of the monetization of Ransomware payments. This trend manifests in two important ways in our data.

First, the median amount of ransom payments to threat actors has been in decline over the past two years, as demonstrated by Figure 45. We want to bring specific attention to the fact that the medians shown in this Figure for 2024 and 2023 are different from the ones we published in the 2025 DBIR. This is because we have added several new data contributors who conduct ransom payments, negotiations and threat actor crypto wallet tracking. Adding these new data sources to our long-standing partnership with the FBI Internet Crime Complaint Center (IC3),[64] who got us started on this whole analysis of ransom impact in 2019, gives us a better view into areas where we previously had lower visibility.

60. They also enjoy long walks on the beach. Away from computers.
61. We wish they would pick better hobbies.
62. Maybe buy them a drink and let them regale you with their tales from the cyber trenches when next you meet.
63. Could they at least have the decency to pick up a couple of tickets and solve some problems while they're in there? No, we didn't think so.
64. ic3.gov

Setting aside the DBIR dataset's version of inside baseball, the takeaway is that our sample size has tripled or quadrupled, depending on the year.[65] In spite of this increase in data volume, the downward trend in the median amount if ransom paid observed from 2023 to 2024 has remained consistent even with the new data sampling,[66] and it has continued from 2024 to 2025. This constitutes relatively strong evidence to support the claim that the downward trend may be real.

Adding to this finding, Figure 46 shows that the percentage of organizations that are not willing to pay the ransomware actors also increased last year – from 65% in 2024 to 69% in 2025. One notable finding from this analysis is that the increase in "Not Paid" outcomes also occurred in cases involving encryption, rather than in data exfiltration events only.

If companies are paying less frequently – even if they have encrypted assets – then the recent attack trend of attempting to inflict the maximum business interruption in order to put a greater time pressure on victims makes even more sense.

Our dataset reveals a market in decline, albeit a slow decline, where there is rampant commoditization and the numerous actors involved are desperately trying to scale to cover their margin compression. The good news here is that the margin compression does not only arise from threat actor competition, but by improved defensive adaptations and increased resilience of the victims. For defenders wondering if their efforts are being successful, the volume reduction of payments to threat actors may be one of those signs of progress that we could measure.

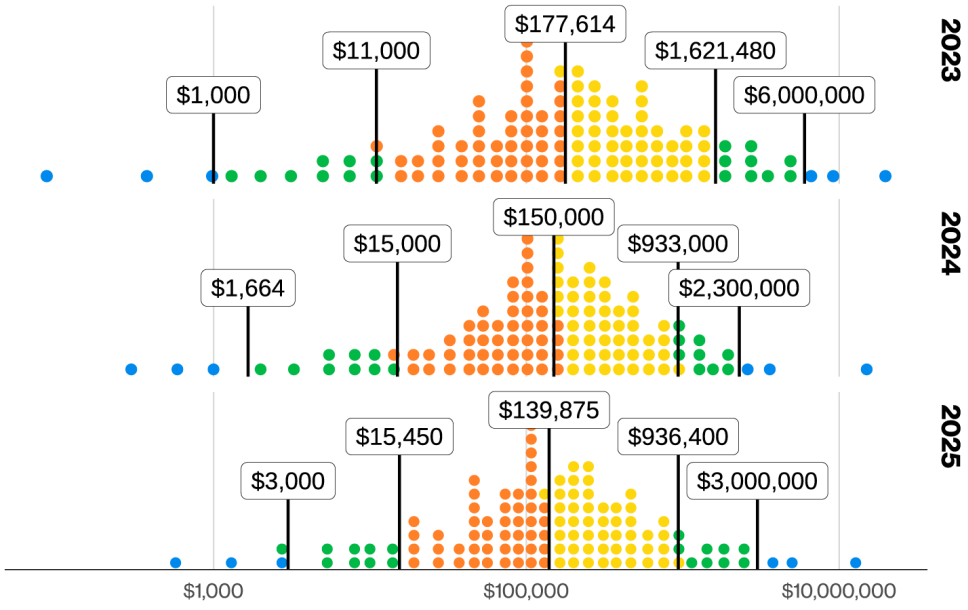

Figure 45. Distribution of loss due to ransom payment in 2025 (n for 2023=1,700 – each dot is 8.50 events) (n for 2024=2,027 – each dot is 10.14 events) (n for 2025=1,494 – each dot is 7.47 events)

Smoke and mirrors in ransomland

Gaining an accurate grasp of the scale of the Ransomware threat has been an ongoing challenge for the industry. The DBIR's generalized approach has largely been to focus on a variety of different sources and combine them so that the biases and limitations can iron themselves out across these sources.

However, the often public nature of a ransomware attack and the way the groups seem to have a constant need to advertise themselves as "reliable Ransomware as a Service providers" complicates matters significantly. There is a growing disconnect between what is being reported and the reality of what has occurred, in no small part due to threat actors reusing old breaches, reposting breaches from other criminal partners and making up breaches out of whole cloth to help increase their notoriety in the criminal world. We're beginning to think that these cybercriminals might not be entirely trustworthy.

To probe this a little further from a different perspective, we cross-referenced data from actor-disclosed ransomware attacks with known actor crypto-wallet payments to estimate how many of the alleged victims actually pay.

After an excruciating amount of time combining, mapping and connecting all the different names, pseudonyms and groupings, the analysis in Figure 47 shows that, based on the reviewed dataset, the median percentage of publicized victims that paid the ransom per ransomware group is about 9%. We cannot do end-to-end tracking of victims due to the nature of the anonymized data, but considering our volume analysis, this sounds like a lot of work for little pay.[67]

65. More data samples is a good thing. Ask any data person.
66. Previous analysis and hypothesis holding steady in a whole new data sample is an even better thing. Ask any data person. Ask a different one than before to increase your sample size!
67. Which, as we all know, contributes to making Jack a dull boy

This result suggests publicized victim records could potentially contain a significant percentage of fabricated entries, as the percentage of payments is under the expected value. This is made worse by the fact that victims that engage and pay the ransom very quickly sometimes are not publicized. As a consequence, the total number of publicized victims per group, which serves as the denominator for these calculations, excludes some of those who actually paid the ransom. This would, of course, skew the results.

This uncertainty complicates even further a problem that was already hard to solve. Policymakers and even individual organizations need accurate information on the impact of ransomware to make better-informed decisions for themselves or their constituents. One of the paths forward involves creating a policy for mandatory reporting of ransomware payments, which has been adopted or is being studied by some countries. The Department of Home Affairs in Australia is one example, and they have provided us with some details of their reporting program on the next page.

Regardless of the percentage of true disclosures, all economic indicators point to the need for threat actors to continue to expand and cast as large a net as possible to get a consistent payday. As this trend of less organizations paying these actors continues, we could end up reaching some sort of saturation point, or perhaps the actors will move toward retirement with a change of heart and a new perspective on their criminal ways.[68]

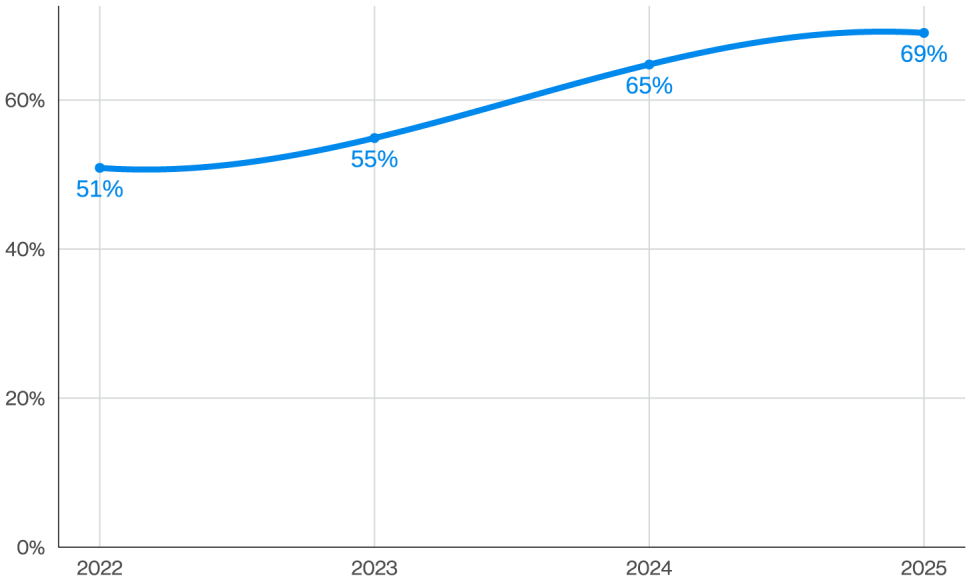

Figure 46. Percentage of ransomware cases where the ransom was not paid (n for 2025 dataset=400)

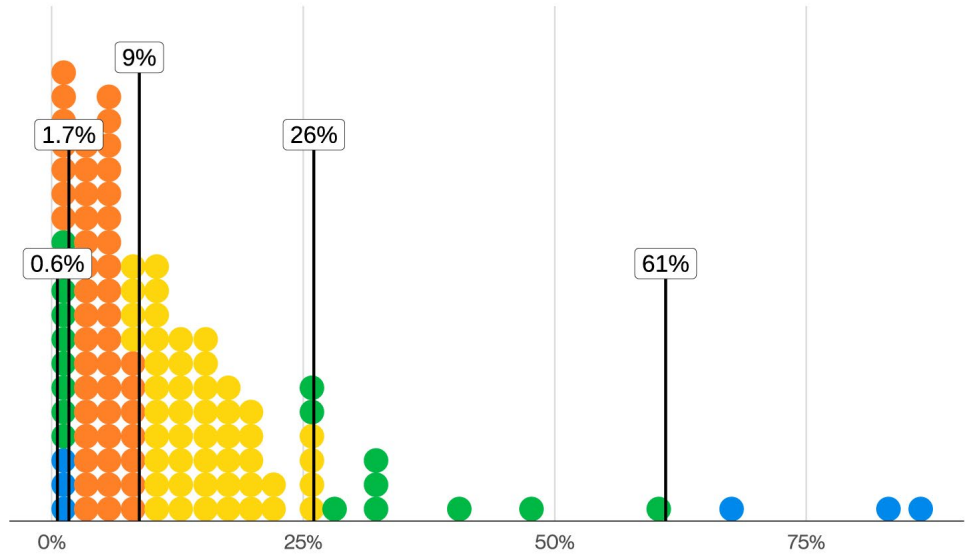

Figure 47. Distribution of ratio of publicized victims to confirmed payments (n=261)

68. And be arrested as soon as they step into a country with extradition agreements with Interpol

Ransomware and Cyber Extortion Reporting Regime

By Australian Department of Home Affairs

The mandatory ransomware and cyber extortion reporting regime commenced on 30 May 2025 under Part 3 of the Cyber Security Act 2024. It requires entities to notify the Australian Government of any payments made in response to ransomware or cyber extortion incidents, with the aim of improving visibility of cyber extortion activity and reducing the profitability of these criminal operations.

The obligation applies to entities with an annual turnover of AUD $3 million or more, as well as operators of critical infrastructure regulated under Part 2B of the Security of Critical Infrastructure Act 2018. The requirement to report is triggered only when a payment has been made, and the report must be lodged within 72 hours of the payment.

Implementation of this regime has commenced with an education first approach. This focuses on helping businesses understand their obligations through guidance materials and engagement activities. From January 2026, we have continued to promote the new regime while moving toward a more mature regulatory posture, maintaining an emphasis on partnership and uplift as organizations adapt to the new requirements.

Infostealer to ransomware pipeline

We are often forced to revisit a topic that we feel we were not able to do justice to for various reasons (never let it be said that the DBIR authors are afraid of beating a dead horse).[69] A case in point is the fact that infostealers as a threat have continued to persist and evolve even in the face of various actions taken by law enforcement. Therefore, we feel justified in retreading some old territory on this topic.

Building on the previous year's approach of analyzing the co-occurrence of infostealer credential leaks and ransomware, we expanded the dataset to include a wider collection of ransomware victims from the two previous years to identify if any of them had infostealers or other types of credential leaks prior to being publicized as a ransomware victim. The analysis found that 27% of ransomware victims had no associated infostealer or credential leak occur within the year. But of the organizations that did, 50% of those ransomware victims had a credential or infostealer event occur within 95 days prior to falling victim to a ransomware attack. Figure 48 shows the distribution of those events, with zero being the day of publication of the victim.

69. Mostly looming deadlines, but more frequently the lack of the right research partner

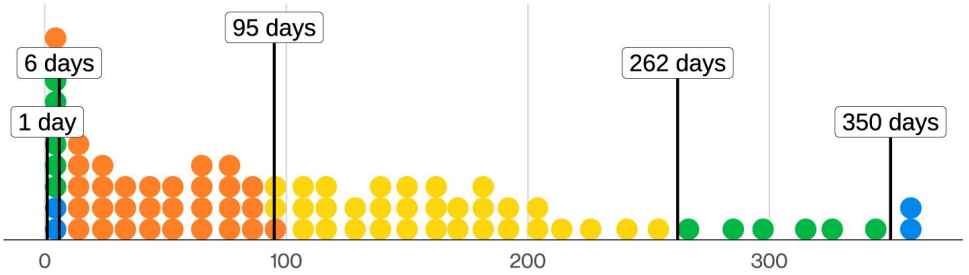

Figure 48. Distribution of days where a credential leakage event occurred prior to ransomware (n=4,395)

In addition, it's not just large organizations that experience credential leakage events in general. Small organizations represented in this dataset experienced a median of seven credential leak events over the course of the year, while larger organizations faced around 20. Although not all of these events mean that an organization is going to experience a ransomware event, they can provide threat actors with an easy entryway that can then be resold to others.

Cost of access

One of the complex elements associated with Ransomware is trying to pin down a pattern across all the various groups, as they rebrand more often than Silicon Valley startups. Individual threat actor groups are known to leverage multiple distinct ransomware toolsets. Adding to the complexity, when trying to determine initial access, some ransomware groups are simply purchasing already compromised access from third parties called Initial Access Brokers (IABs).

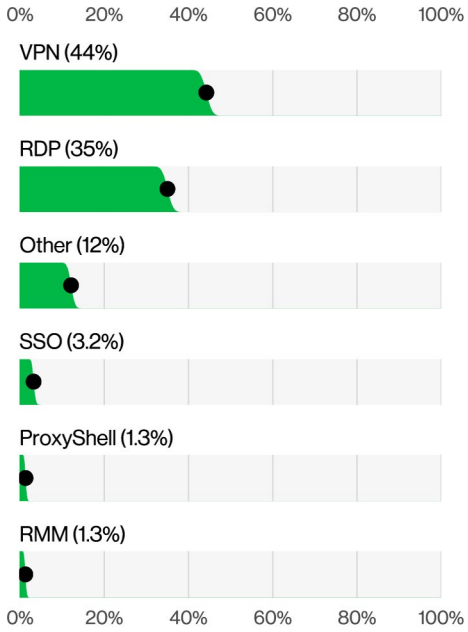

Figure 49. Percentage of select known connection types by IAB offerings (n=876)

By outsourcing the target acquisition, ransomware operators are able to focus on their skillset of lateral movement, privilege escalation, deploying their ransomware payload and, well, potentially getting paid.

By examining offers made by these IABs, we can get a general understanding of the types of access they are offering and the associated prices. There is large diversity in terms of the different access vectors to the environments, everything from more traditional VPN access to specific application servers.[70] Figure 49 has a breakout of these types of connections across the last four years. Without much surprise, we see that 44% of the connection types are VPN, followed closely by some type of remote desktop application (e.g., RDP, RDPweb, VNC). Also interesting is the inclusion of ProxyShell/ProxyLogon access, considering the majority of these credential offers were found to have occurred two to three years after the vulnerability was disclosed, once again reiterating the value of the long-tail of our incomplete patching and mitigation process to attackers.

70. Wanting access to more than a single project management tool instance is clearly a hallmark of the criminally insane.

When it comes down to the actual value of these types of access, Figure 50 shows that regular, non-privileged accounts were typically worth around $700. Administrative accounts, however, were worth almost double that, valued at around $1,300 per account. The DBIR's function is not to provide recommendations or suggestions to threat actors, but we were expecting a higher median price, to be honest.

Median analysis aside, the contrast between a basic account and a privileged one on the extreme end of the distribution shows how much attackers value these high-privileged accounts, as they allow them to bypass one of the key steps required for them to achieve their missions: privilege escalation. There is a section later in the report dedicated to a deep dive regarding our observations on privilege escalation.

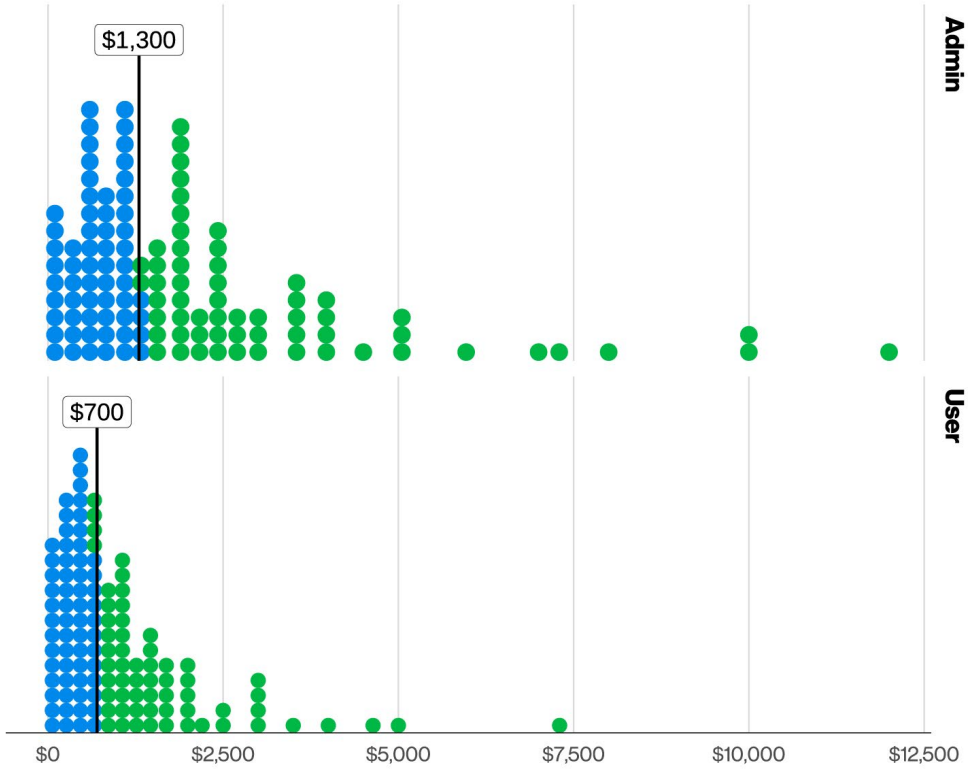

Figure 50. Cost of IAB offerings per access level (n for Admin=453) (n for User=1,008)

CIS Controls for consideration

Protecting devices

Secure Configuration of Enterprise Assets and Software [4]
- Establish and Maintain a Secure Configuration Process [4.1]
- Establish and Maintain a Secure Configuration Process for Network Infrastructure [4.2]
- Implement and Manage a Firewall on Servers [4.4]
- Implement and Manage a Firewall on End-User Devices [4.5]

Email and Web Browser Protections [9]
- Use DNS Filtering Services [9.2]

Malware Defenses [10]
- Deploy and Maintain Anti-Malware Software [10.1]
- Configure Automatic Anti-Malware Signature Updates [10.2]

Continuous Vulnerability Management [7]
- Establish and Maintain a Vulnerability Management Process [7.1]
- Establish and Maintain a Remediation Process [7.2]

Data Recovery [11]
- Establish and Maintain a Data Recovery Process [11.1]
- Perform Automated Backups [11.2]
- Protect Recovery Data [11.3]
- Establish and Maintain an Isolated Instance of Recovery Data [11.4]

Protecting accounts

Account Management [5]
- Establish and Maintain an Inventory of Accounts [5.1]
- Disable Dormant Accounts [5.3]

Access Control Management [6]
- Establish an Access Granting/ Revoking Process [6.1, 6.2]
- Require MFA for Externally-Exposed Applications [6.3]
- Require MFA for Remote Network Access [6.4]

Security awareness programs

Security Awareness and Skills Training [14]

Social Engineering

Summary

Threat actors continue to largely leverage email-based phishing attacks to compromise organizations; however, these attacks are getting more complex as attackers are targeting mobile devices and other unconventional vectors to reach victims.

What is the same?

Email is still the preferred attack vector for the majority of Social Engineering breaches.

Frequency	5,302 incidents, 3,814 with confirmed data disclosure
Threat actors	External (100%) (breaches)
Actor motives	Financial (86%), Espionage (25%) (breaches)
Data compromised	Other (56%), Internal (51%), Credentials (39%), Secrets (31%) (breaches)

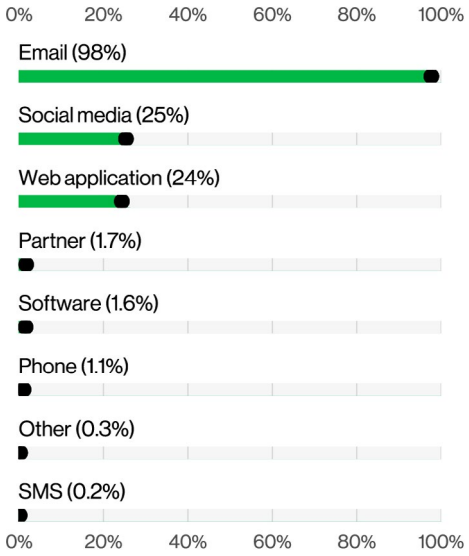

Figure 51. Top Social vectors in Social Engineering breaches (n=3,777)

Social Engineering and phishing via email have been more or less considered synonymous in the defender's playbooks for some time now. For those of us who have been around long enough, we still have memories of the early days of AOL and our first email accounts, along with the desperate pleas from rich princes in foreign lands in need of help from someone trustworthy whose only qualifications were owning an email address.[71]

While some things have changed over the course of the last few decades, some have also stayed the same, with Social Engineering continuing to be one of the most common types of attack resulting in breaches since 2018. The Social Engineering pattern shows up as our third most common breach pattern, representing 16% of breaches. But like everything else in cybersecurity, change and evolution are inevitable.

The social quo

This pattern is mainly focused on the attacks that leverage deception to accomplish some specific objective, be it deploying malware, collecting credentials or transferring money to the actor's bank account. The common element in this pattern is that individuals are targeted in the attack.

Typically, in these types of attacks, the first step against the victim is the Social action, as it is what provides the attacker access to an organization's environment. Figures 51 and 52, respectively, show the Social action varieties (types of attacks) and vectors (the distribution methodology). As expected, we see Phishing and Email as the primary methods used.

When we're looking at data from email security gateways, we see similar breakdowns (Figure 54), with 80% of the attacks blocked being plain phishing, 10% being emails with malware, 5% being attempts of getting the victim to call back to the attacker and 3% consisting of Business Email Compromise-style attacks such as attackers trying to get the victim to update an existing bank account (or register a new one) ahead of a wire transfer, usually by pretending to be that user by leveraging the historical email chain. Those last two are actually good examples of Phishing followed by another social action, Pretexting, due to the synchronous nature of the attack. If there is someone on the other side of the proverbial line interacting with you to do something you shouldn't, that's Pretexting.

71. And a bank account

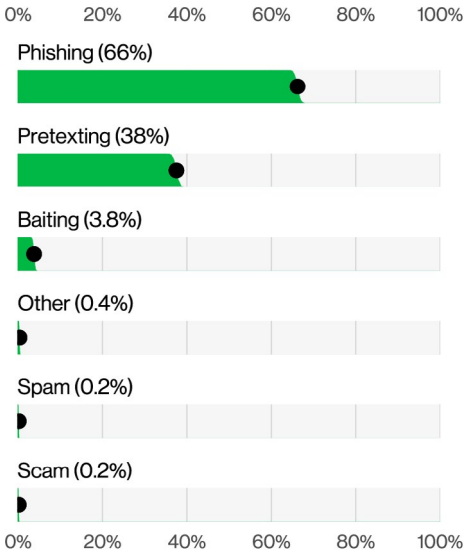

Figure 52. Top Social actions in Social Engineering incidents (n=5,217)

When examining how much these mobile devices are getting phished via SMS, as Figure 53 shows, we find that the median amount per year of those SMS-based phishing campaigns targeting mobile devices in large organizations is 48 (12 for smaller organizations). While these numbers don't look eye-poppingly large,[72] they do demonstrate an existing and present threat that can bypass traditional approaches to phishing mitigation by reaching directly to the users through their mobile devices.

It's important to point out that those detections were only visible because those are managed devices – either corporate owned or with some form of mobile device management software. If your employees are using purely unmanaged personal devices to perform organizational duties, this can represent a risky gap in your visibility.

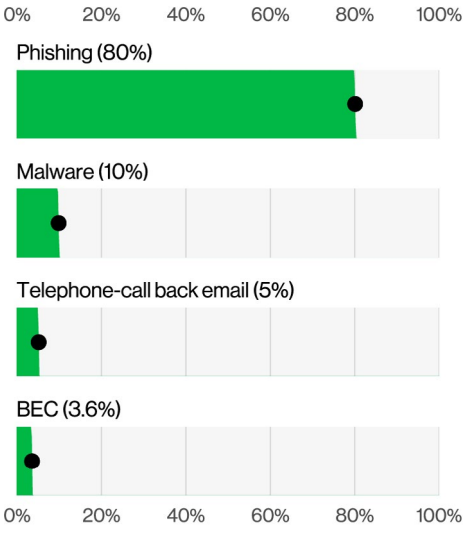

Figure 54. Median percentage of Email attack types by month (n=3,709,045,972)

Multipronged social approach

Much like how advertisers try to grab our attention through a variety of different methods – web applications, emails, text messages and airplane banners at the beach – so do attackers. In fact, 41% of our Social Engineering breaches involve social vectors other than just Email. In this section, we'll look at the types of attacks that go beyond just the Email vector and explore the other methods that attackers are leveraging to target our employees and the implications to our defenses.

When we're looking at our incident data, we see that about a fourth of our social action vectors come from either Social media or Phones, which really represents the widening of the net that attackers are leveraging to snare our users.

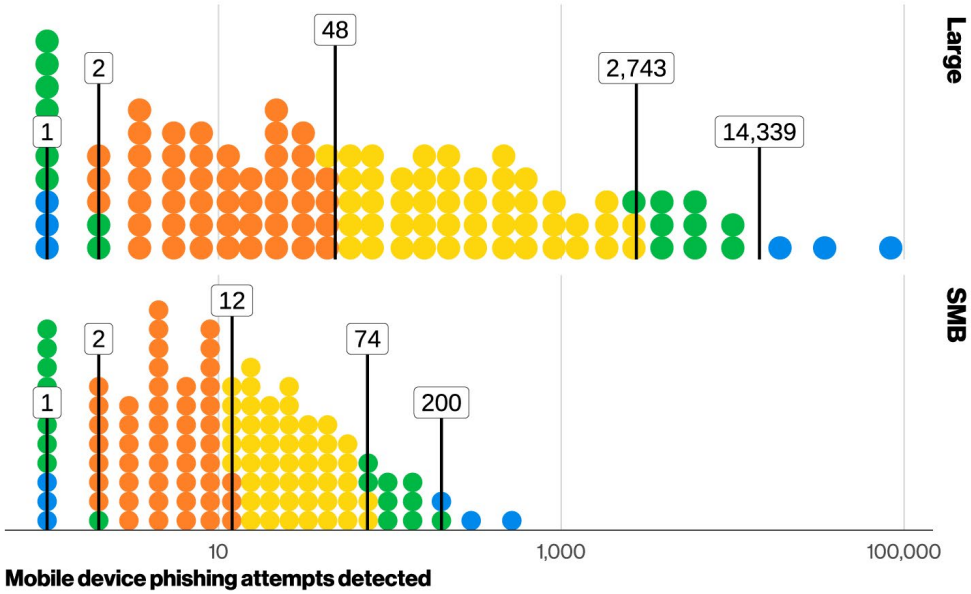

Figure 53. Annual number of mobile device Phishing attempts detected by org size (n=12,363 – each dot is 103.025 attacks)

72. We'd argue that being targeted by such a campaign every eight calendar days in the year is worth looking into.

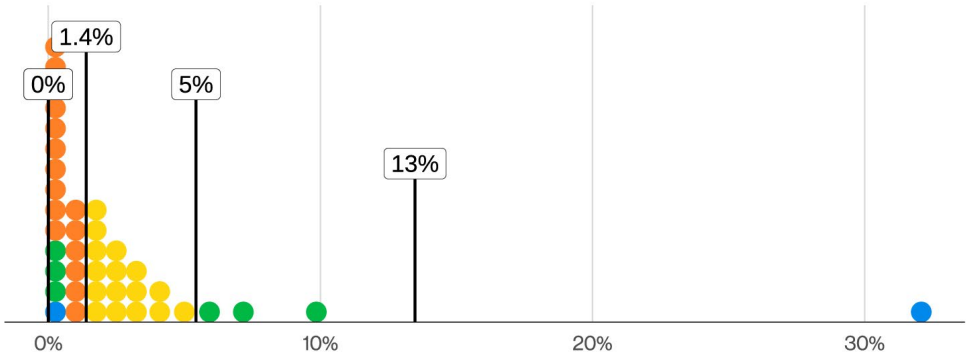

Figure 55. Distribution of success rate of Email vector-simulated social attack campaigns (n=8,395 – each dot is 209.88 campaigns)

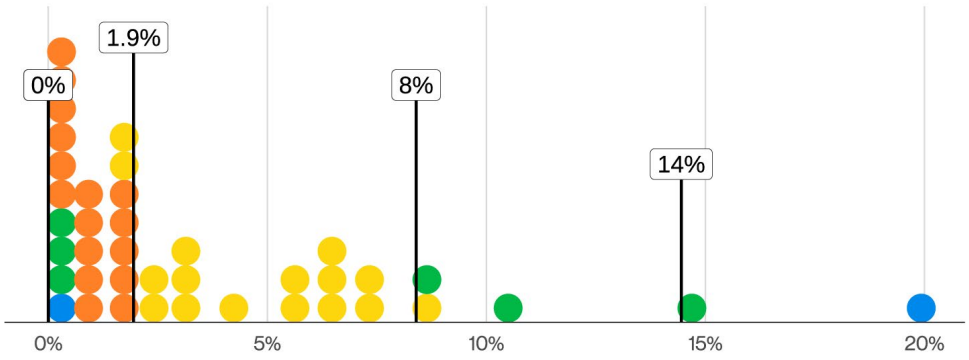

Figure 56. Distribution of success rate of non-Email vector-simulated social attack campaigns (n=35 – each dot is 0.88 campaigns)

Aside from receiving sneaky links for users to click on, phones are also the perfect tool for voice phishing, which is just another way of saying Social Engineering over a phone call. Now, if you have been following along, this is where we put our VERIS hats back on and reinforce that those are classified as Pretexting in our dataset, which has seen a substantial increase in our initial access vector analysis from last year. Regardless of the terminology, various attackers have been leveraging these means, by impersonating help desk agents or users needing a password reset, with moderate levels of success.

The bottom line here is that social attacks using phone-centric vectors – text messages, voice or the previously mentioned callback-focused emails – are more successful in our dataset than using the traditional Email vector defenders are used to. While the median click rate of email phishing simulation campaigns is 1.4%, we see the median rate of simulations on phone-centric methods is closer to 2%, as demonstrated in Figures 55 and 56. That is an increase of 40% in the median click rate between those vectors.

The caveat on this analysis should become apparent when you look at the sample sizes. Phishing simulation is a well-established business practice, and we have a handful of leading companies in this space as data contributors. On the other hand, we struggled to find companies doing simulations of voice- and text message-based campaigns, which leads to this small-ish sample size. We hope that, for the 2027 DBIR, we will be able to collect more data as additional companies that offer these kinds of services want to participate in this research.

Regardless, the result suggests that different strategies of training and simulation are needed to mitigate the risk of those "new" vectors. The more involved Pretexting attacks on the rise are also of a different nature to the "send-message-and-hope-for-a-click" Phishing attacks, as they are tailored to appeal to the nature of the employees being targeted.[73] The takeaway reflection here is this: How are users taught to detect these unconventional social engineering attacks? Could you detect someone impersonating your help desk, and by what means can your users be reached on your devices?

73. What do you mean the help desk is not supposed to help someone, in some cases?

In the trenches

One common breach structure we reviewed this year involves combining multiple elements to make for a convincing scenario to request and obtain access to internal systems.

1. Attackers will create some fake type of IT emergency for users, for example, by signing them up for various spam services, which results in the users getting bombarded with suspicious emails.
2. Fortunately for the users (and unfortunately for the organizations), a "helpful" individual sends an external chat request via Microsoft Teams[74] claiming to be from the help desk and offers to help if the users just share access to their desktop systems.
3. With this access, the attacker is able to conduct the attack from the victims' user devices while "troubleshooting" for the users.

For defenders, detecting those types of attacks can be very challenging, as there's no malicious code being executed, the access point is via a messaging app that oftentimes by default will allow external entities to send in connection requests and the remote access usually leverages tools built into the operating system and approved by the organization.

The (click)fix is in!

At this point, you have probably sworn off of ever checking your email or picking up your phone again. Don't worry, threat actors have a Social action just for you! Welcome to the world of Baiting, where attackers set up realistic webpages, create online adverts, or compromise existing websites to get users to download their malware via crafty webpages posing as downloads of legitimate tools and software. These can show up in a few different varieties, but the main types that we've encountered this year are the search engine optimization (SEO) abuse downloads and the ClickFix attack.

The ClickFix attacks are a newer twist on an old tactic, where the malicious webpages present themselves as a CAPTCHA, which in itself is pretty familiar to any modern-day netizen.

74. According to publicly available incident reports, Teams has been a very common vector along with Quick Assist, but this does not preclude the involvement of other messaging systems.

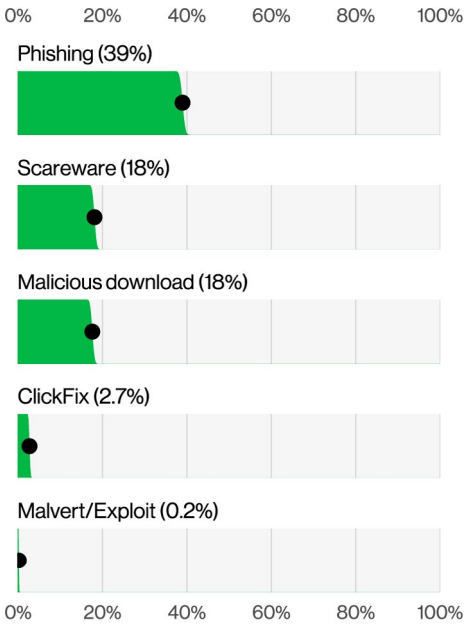

Figure 57. Attack types detected at the browser (n=8,498)

However, there is a twist! This CAPTCHA isn't about solving some indecipherable scribbles or pressing a button but is prompting the user to open up a terminal window and paste a command into that terminal.[75] Naturally, the command is some malicious payload that promptly downloads actual malware onto the system.

While this may seem obvious in a controlled setting, these attacks are not reserved exclusively for use against non-technically savvy users. In practice, attackers skillfully combine technical instructions with Social Engineering to instill a sense of both urgency and distraction. These psychological pressures are designed to bypass users' typical caution, leading them to execute commands (in this case, press Ctrl-Alt-R and then Ctrl-V) they would otherwise recognize as a threat.

When we examine attacks blocked at the browser level (Figure 57), we find that these ClickFix attacks only represent about 2.7% of attacks, still being overshadowed by the traditional phishing pages and malicious downloads. These attacks are interesting from the novelty perspective, but also for the implications for defenders in how they can detect and mitigate users copying and executing commands from a webpage.

Tactical Espionage social action

When it comes to Social Engineering, it's clear that the APT Crowd is living up to the title of persistent. About 25% of our Espionage-motivated incidents involve Social Engineering actions as part of the attack path. Where these actions play out is through long-term interactions and the development of rapport with the targets.

One such tactic these attackers leverage is by weaponizing a job hiring process, such as attempting to recruit employees of the target organization as means of getting internal information. Alternatively, they may try getting the user to download and troubleshoot a git repository that happens to have malware embedded in it. As our analysis on bring your own device (BYOD) and infostealers in the 2025 DBIR shows, even on personal devices, it was relatively common for employees to have corporate account credentials or information that is ripe to be compromised. When it comes down to helping employees protect themselves in situations like these, the focus should also be on promoting the same discerning sensibility at home in their personal lives as they do at work.

75. Be very careful which link you click in your web search if you are trying to install legitimate open-source software.

CIS Controls for consideration

Protect accounts

Account Management [5]
- Establish and Maintain an Inventory of Accounts [5.1]
- Disable Dormant Accounts [5.3]

Access Control Management [6]
- Establish an Access Granting/ Revoking Process [6.1, 6.2]
- Require MFA for Externally-Exposed Applications [6.3]
- Require MFA for Remote Network Access [6.4]

Security awareness programs

Security Awareness and Skills Training [14]

Managing incident response

Incident Response Management [17]
- Designate Personnel to Manage Incident Handling [17.1]
- Establish and Maintain Contact Information for Reporting Security Incidents [17.2]
- Establish and Maintain an Enterprise Process for Reporting Incidents [17.3]

Basic Web Application Attacks

Summary

Basic Web Application Attacks remain widespread and are typically driven by stolen credentials and unpatched vulnerabilities. While often low in sophistication, they are highly effective and frequently lead to credential theft, internal data exposure and further compromise of systems.

What is the same?

The Use of stolen creds continues its historic run in this pattern, showing the more things change, the more they remain the same.

Frequency	3,217 incidents, 2,281 with confirmed data disclosure
Threat actors	External (100%) (breaches)
Actor motives	Financial (74%), Espionage (23%), Ideology (3%) (breaches)
Data compromised	Credentials (52%), Internal (48%), Other (33%), Secrets (15%) (breaches)

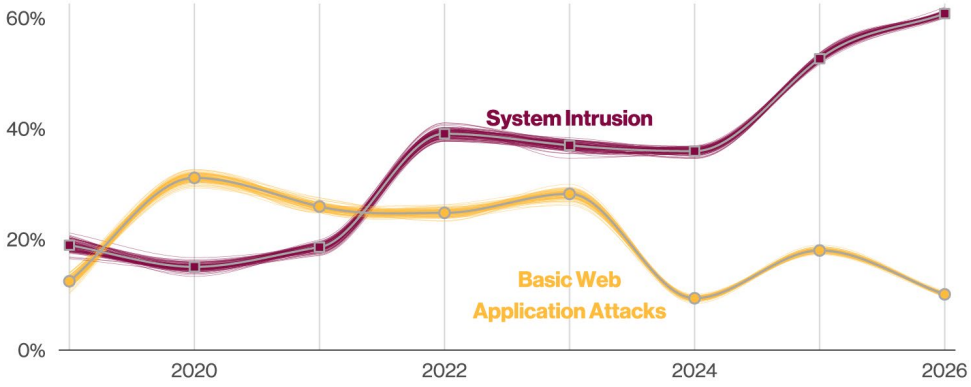

Figure 58. System Intrusion and Basic Web Application Attacks breaches over time (n for 2026 dataset=22,624)

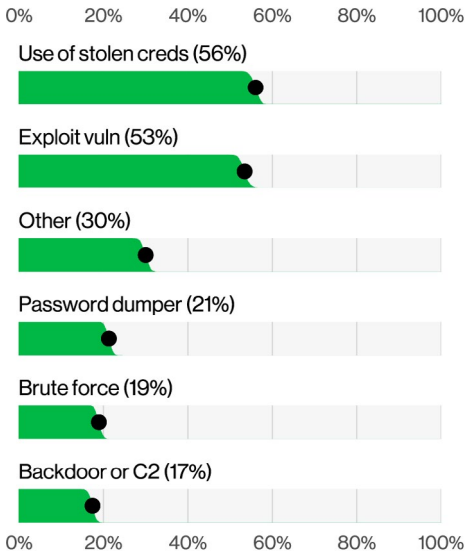

Figure 59. Top Action varieties in Basic Web Application Attacks breaches (n=2,193)

The Basic Web Applications Attacks pattern is, as the name suggests, relatively straightforward. Rather than the intricate, movie-style heists orchestrated by brilliant criminal masterminds, it's closer akin to "oh, they left the door open" on the spectrum of attack complexity. Figure 58 shows us that it has been trending down in favor of the more complicated System Intrusion. This trend suggests that defenders may be successfully raising the bar, thereby making attackers work harder (or attackers are leveraging their initial access to achieve more complex objectives).

Once again, the Use of stolen creds is the top action in this pattern (Figure 59). How attackers obtain these credentials is often unknown. They may be coming from phishing, infostealers or from data exposed in prior breaches and later packaged and sold on the dark web. However they acquire them, they are making good use of them in this pattern.

We also saw a rise of the Exploit vuln action this year due to several large cases involving software vulnerabilities that were left unpatched and later exploited – either in an organization's own infrastructure or that of a partner.

Also making a debut this year is Password dumper. It is not uncommon for the perpetrators of this type of attack to harvest additional credentials for use in further attacks or to simply monetize them. It is closely followed by Brute force,[76] which is an excellent way of breaking the passwords when you cannot steal them.

The number of incidents and breaches classified under this pattern rose significantly compared to 2025's report. This could mean the number of incidents increased, but it could also mean our partners had improved visibility into the low-effort attacks that typically fall into this pattern. As always, it is difficult to attribute such changes to a single cause when looking at our data.

We observed an increase in financially motivated actor breaches compared to last year, while Espionage-related breaches declined. Incidents driven by Ideology, however, remained fairly consistent.

CIS Controls for consideration

Mitigation efforts against stolen credentials

Account Management [5]
- Establish and Maintain an Inventory of Accounts [5.1]
- Disable Dormant Accounts [5.3]

Access Control Management [6]
- Establish an Access Granting/ Revoking Process [6.1, 6.2]
- Require MFA for Externally-Exposed Applications [6.3]
- Require MFA for Remote Network Access [6.4]

Mitigation efforts against vulnerability exploitation

Continuous Vulnerability Management [7]
- Establish and Maintain a Vulnerability Management Process [7.1]
- Establish and Maintain a Remediation Process [7.2]
- Perform Automated Operating System Patch Management [7.3]
- Perform Automated Application Patch Management [7.4]

76. They are the Brute Squad!

Miscellaneous Errors

Summary

Employee mistakes – especially Misdelivery and Misconfiguration – remain a persistent and significant cause of data breaches, often exposing personal data. Better access controls, monitoring and safeguards against common human errors are critical to reducing these incidents.

What is the same?

Errors were the cause of breaches slightly more often than last year. Misconfiguration and Misdelivery continue to be the primary issues organizations are experiencing in terms of the kinds of errors their employees are making.

Frequency	1,757 incidents, 1,750 with confirmed data disclosure
Threat actors	Internal (100%) (breaches)
Data compromised	Personal (98%), Internal (16%), Other (8%), Bank (7%) (breaches)

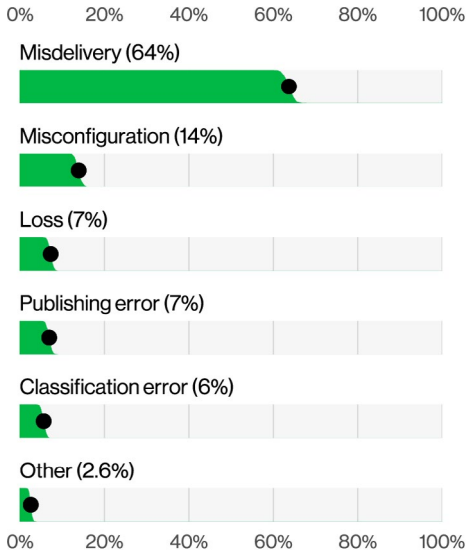

Figure 60. Top Action varieties in Miscellaneous Errors breaches (n=1,719)

Better to ask forgiveness?

There is an old adage that it is better to ask forgiveness than permission, but it doesn't really age well when it comes to the kinds of mistakes we repeatedly observe in this pattern. Some industries (we're looking at you, Healthcare) have had this pattern in their top three for as long as we have been tracking it. Surely the organizations (and the people whose data has been compromised) are rather tired of people doing things in the name of expediency (or convenience) or other reasons that result in a breach. Forgiveness is generally not a concept recognized in most data breach disclosure laws.

We witnessed some movement in the top three Errors for this pattern this year. Figure 60 illustrates that Misdelivery, Misconfiguration and Loss were the most common, with Publishing errors not far behind. Last year, Publishing errors were in third place. For those new to the report, Publishing errors occur when private data is mistakenly placed on a public-facing server. These are typically discovered when search engines index the content and customers searching their own name receive an unpleasant surprise.

Misconfiguration errors, by contrast, occur when someone deploys a data store on the internet without applying appropriate access controls. These are most commonly discovered by security researchers, who then attempt to make a notification if they can determine whose data it is. What we don't know is how often other, less civic-minded people have encountered the same data, made a copy and quietly slipped away. The fact that Misconfiguration remains among the top errors over time is rather concerning.

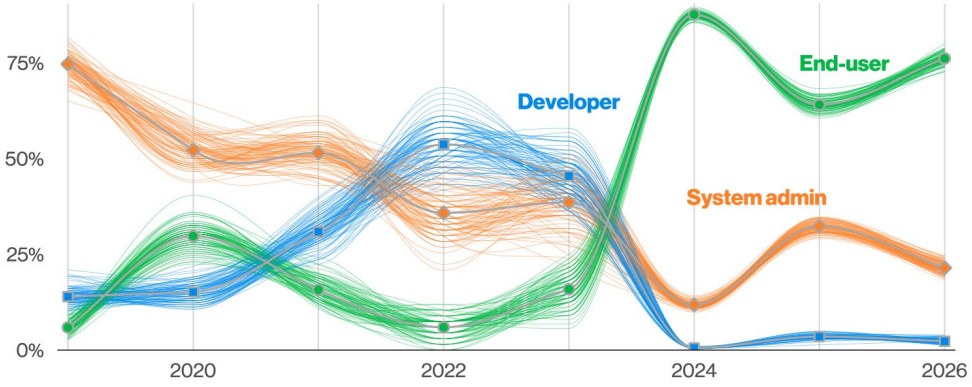

Figure 61. Select Actor varieties over time in Miscellaneous Errors breaches (n for 2026 dataset=861)

Misdelivery errors occur when data is delivered to the wrong recipient. While we still see this happening with documents, particularly for those organizations that do large mass mailings, it also occurs frequently in electronic form, as well. It is very easy to make this type of mistake—does your organization have a way of intercepting these kinds of errors before there is a full-blown data breach?

One useful step would be to identify who most often causes these Errors in your organization. We can see in Figure 61 that Developers, System admins and End-users top our list. Certainly End-users will be more numerous at any organization, but Developers and System admins tend to have access to higher volumes of data and privileged access, which increases the potential impact of this type of mistake.

Organizations that want stronger access controls in place may want to limit who has access to their most critical data and place safeguards and procedures accordingly.

CIS Controls for consideration

Control data

Data Protection [3]
- Establish and Maintain a Data Management Process [3.1]
- Establish and Maintain a Data Inventory [3.2]
- Configure Data Access Control Lists [3.3]
- Enforce Data Retention [3.4]
- Securely Dispose of Data [3.5]
- Segment Data Processing and Storage Based on Sensitivity [3.12]
- Deploy a Data Loss Prevention Solution [3.13]

Secure infrastructure

Continuous Vulnerability Management [7]
- Perform Automated Vulnerability Scans of Externally-Exposed Enterprise Assets [7.6]

Application Software Security [16]
- Use Standard Hardening Configuration Templates for Application Infrastructure [16.7]
- Apply Secure Design Principles in Application Architectures [16.10]

Train employees

Security Awareness and Skills Training [14]
- Train Workforce on Data Handling Best Practices [14.4]
- Train Workforce Members on Causes of Unintentional Data Exposure [14.5]

Application Software Security [16]
- Train Developers in Application Security Concepts and Secure Coding [16.9]

Privilege Misuse

Summary

Intentional insider misuse is less common than External attacks, but privileged access, convenience-driven policy violations, AI data leakage and fraudulent employee activities are growing insider-related risks that organizations must monitor closely.

What is the same?

Financially motivated insiders continue to steal data to benefit them down the road. Whether it is taking it to another employer or starting a competing business, these are notoriously difficult to detect.

Frequency	1,141 incidents, 766 with confirmed data disclosure
Threat actors	Internal (100%), External (1%), Multiple (1%) (breaches)
Actor motives	Convenience (60%), Financial (33%), Espionage (4%), Grudge (4%), Fun (2%), Other (2%) (breaches)
Data compromised	Personal (60%), Other (35%), Secrets (27%), Internal (25%) (breaches)

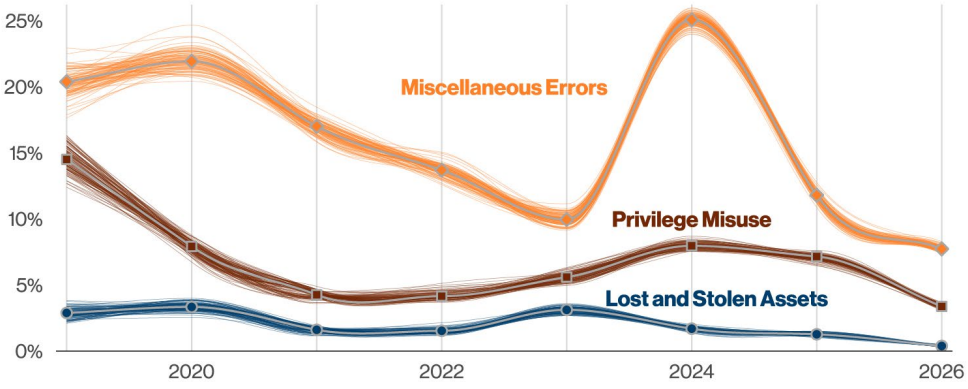

Figure 62. Select patterns over time in Privilege Misuse breaches (n for 2026 dataset=22,624)

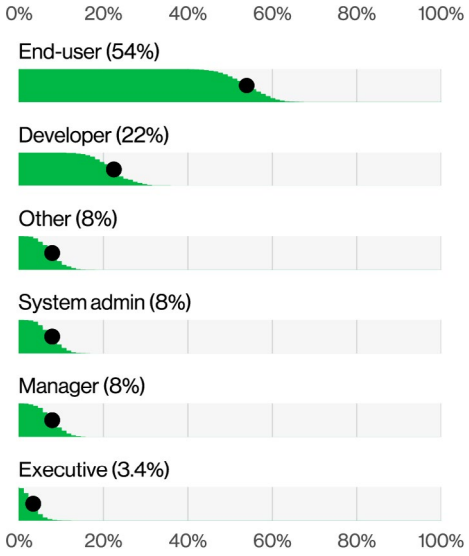

Figure 63. Top Actor varieties in Privilege Misuse breaches (n=89)

As Figure 62 illustrates, the Privilege Misuse pattern has never been a dominant driver of data breaches. Organizations generally face far more risk from External actors than from their own employees. This doesn't mean insider malice is insignificant, and when it occurs, the impact is rarely trivial. However, your own people are typically much more likely to lose an asset or make a mistake than to intentionally abuse their access.

Misuse breaches have declined over the past two years, after peaking at nearly 8% of breaches in the 2024 report. In this year's report, they account for less than half that figure at just under 4% of breaches. By comparison, the Miscellaneous Errors pattern represents nearly 9% of breaches this year.

Some of this can be attributed to changes in our contributor base and the degree to which they can detect and report malicious insider breaches. Even so, the Privilege Misuse pattern has remained a relatively minor player overall.

The trust gap

Now that we have established that these types of breaches occur less frequently than some might think, when they do occur, their impact can be exceptionally damaging to organizations. This leads us to the obvious question: Who do we see perpetrating serious violations against their employers? In most cases (54%), it is your average End-user (Figure 63). However, we also see Developers and System admins being involved – which is especially concerning given that these two account types tend to have higher levels of privileged access. Finally, Managers are also sometimes responsible for compromising data, and most frequently, these are financially motivated and are intended to benefit them at future employers.

This year, we observed an increase in the number of these breaches that had Convenience as a motive. A common example is when an employee wants to work from home and emails company data to a personal account. That would be a case where the person wasn't specifically acting with malicious intent – they wanted to keep working – but they clearly violated the data-handling policies of their employer and caused a data breach. We will be watching closely to see whether this type of Convenience-related breach continues to become a trend.

Coercion of high-risk employees

Over the years, many of our data contributors in DLP and behavioral engineering have shared a very clear datapoint: Policy violations are rarely evenly distributed. There are always "troublemakers" inside organizations, in the sense that many policy violations are often concentrated in a small group of individuals.[77] Even when unintentional and accidental, there remains a small subset of individuals who account for a disproportionately high level of risk.

However, one under-explored topic around very high-risk individuals is their susceptibility to coercion-based collaboration with External actors. Coercion tactics often require a blackmail target, and recent technological developments, along with the availability of personal information that might make for newsworthy topics, have made it much easier to obtain leverage against those high-risk individuals.

More to the point, research from one of our newest data contributors[78] shows that across their installation base of almost 270,000 enterprise work computers, approximately 1 in 500 employees accessed high-risk compromising materials on an enterprise device.

This includes content related to extremism, promotion of bodily harm or exploitative materials not appropriate for a workplace environment. The existence of such compromising materials has often been associated with an increase of how susceptible to coercion an employee can be. If those employees have any sort of privileged access in your organization, their trusted access could become externally leverageable.

Military and intelligence agencies have long incorporated coercion-based insider risk mitigation procedures in their environments. Given the documented increase we have seen of state-sponsored activity targeting the private sector, it is definitely worth considering when iterating your organization's insider risk program.

77. A Misuse Pareto, if you will
78. NetClean has more details on their insights section at netclean.com/knowledge/insights-and-data.

Insider breaches can be very difficult to detect in real time because they are frequently using access legitimately granted to them to perform their duties. The question is whether your controls can identify when these actors go rogue. If you are unsure what this kind of breach might look like, perhaps a good starting point would be to begin with the offboarding process for those people with access to the most sensitive data in your organization. For example, triggering a simple activity review when someone on that list resigns (or is let go) may bring to light a breach that would have otherwise remained undetected until that data is misused later. This simple control has caught numerous breaches in our dataset.

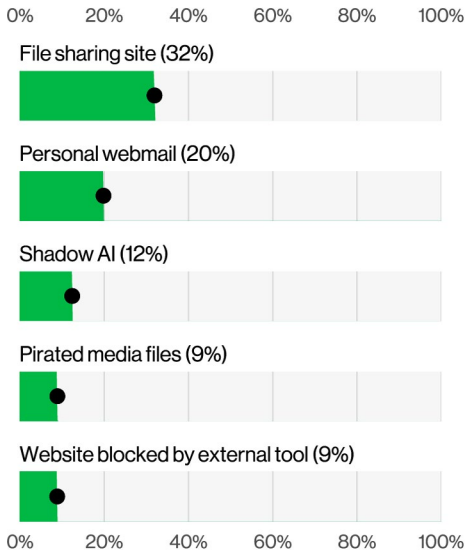

Figure 64. Top non-malicious insider untrusted DLP event targets (n=4,280,149)

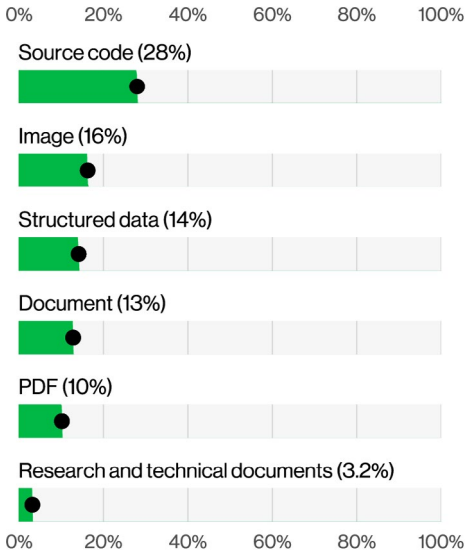

Figure 65. Select data types in untrusted DLP events targeting generative AI tools (n=858,440)

DataGPT: gone, pilfered or transferred

Just as we can't go through a report without mentioning vulnerabilities and credentials at least 56 times, we're also going to be banging the drums around GenAI. Specifically, this section focuses on the possible exposure of our most precious commodity, our data. The last few years could perhaps be best described as having a blistering rate of adoption of this new technology, as now 45% of employees are considered regular users[79] of AI on their corporate devices, up from 15% that we reported last year. However, while the user adoption is growing, it's unclear if security is either playing whack-a-bot or driving and informing these decisions as part of a comprehensive AI usage policy.

When it comes to the provenance of the accounts being used on AI platforms, we find a similar type of breakout as last year, with 67% percent of users using non-corporate accounts on their corporate devices to access AI platforms.[80] These types of unaccounted AI systems that contain corporate data are sometimes referred to as "Shadow AI." Much like Shadow IT,[81] these systems exist outside of the control of the organizations and can represent a significant risk for data leakage. This issue has become so prevalent this year that it is now the third most common non-malicious insider action detected in our DLP service datasets in 2025 (Figure 64), a fourfold increase from last year.

And it's not just bots we have to be careful of. Even our portal to the internet, the browser, is becoming increasingly more integrated with AI. We found this year that the average company had more than 15% of users with unauthorized AI extensions installed on their browsers. One of the main functions of these AI plugins is often to collect and retain information about what the user is browsing for context, so if your corporate users are browsing your internal sites, some of your non-public data might be getting vacuumed up

When it comes to where our data goes and whose probabilistic model gets to pontificate on it, we found that the most common data submitted to external AI models was source code, by a large margin (Figure 65), followed by images and other types of structured data.

79. In our dataset, that means access to an AI platform at least once every 15 days.
80. Last year, that percentage was 72%.
81. And unlike Shadow the Hedgehog, even though he also ended up outside of the control of his organization

In 3.2% of DLP events, we even found research and technical documentation being uploaded to untrusted and unauthorized AI systems, possibly leaking key internal research. As if the source code part was not enough, you now have potential intellectual property walking out the door.

Considering the amount of data that these models consume, process and log, how comfortable should we be with uploading our secret recipes and key intellectual property data to these unauthorized third parties? Especially considering the increasing number of systems, servers and hands that reside between your user and the actual model.[82] Just because there's a new toy to play with doesn't mean we should ignore decades of data governance and third-party risk management practices. Roko's basilisk[83] is not real and it cannot hurt you!

CIS Controls for consideration

Manage access

Secure Configuration of Enterprise Assets and Software [4]
- Establish and Maintain a Secure Configuration Process [4.1]
- Manage Default Accounts on Enterprise Assets and Software [4.7]

Account Management [5]
- Disable Dormant Accounts [5.3]
- Restrict Administrator Privileges to Dedicated Administrator Accounts [5.4]

Access Control Management [6]
- Establish an Access Granting Process [6.1]
- Establish an Access Revoking Process [6.2]

82. If it wasn't clear, this is a rhetorical question.
83. en.wikipedia.org/wiki/Roko%27s_basilisk

Denial of Service

Summary

DDoS extremes continue to increase as organizations face erratic burst attacks year-round, with the median breached entity contending with 17 distinct attacks throughout the year.

Frequency	5,514 incidents, 3 with confirmed data disclosure
Threat actors	Internal (100%) (breaches)

What is the same?

DDoS attacks continue to be one of the top incidents targeting a wide variety of different industries.

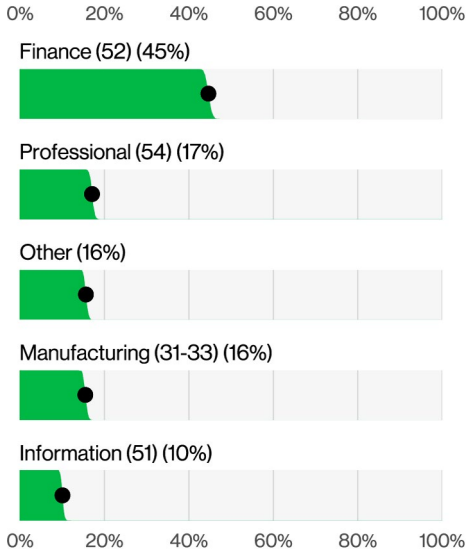

Figure 66. Top victim industries in DDoS incidents (n=5,513)

As our world continues its movement toward interconnectedness (and possibly the singularity), the systems and networks that connect us become increasingly more important every day. While many of the newsworthy outages this year were due to DoS accidents, rather than DoS attacks, it doesn't mean that these threats haven't been a consistent and thorny issue for many organizations.

Down but not out

The first step into our analysis is simply to understand who is being impacted by DDoS (or who knows that they were targeted), as Figure 66 shows. What we see are the same industries that we've seen for the last couple of years, just with some shifting in terms of rank ordering. While in the 2025 report there was a relatively large spike in DDoS events targeting the government and the wider Public Sector, possibly due to hacktivism, they've settled back down to their normal spot near the bottom. On the top, we continue to see that Finance, Professional Services and Manufacturing are the typical main targets, as they have been since 2022.

Extreme growth

Channeling our inner 90s kids, rocking our skateboards and inexplicably calling everyone "Brah," we wanted to look a bit into the extremes of the DDoS dataset. But to first understand the extremes (what we're really talking about are the outliers), we need to understand the things that are essentially the baseline in this dataset. In comparison to previous years, there has been some light jockeying in terms of the typical size, volume and duration of an attack. The median size of an attack has fluctuated between 4.2 Gbps and 6 Gbps over the last few years, with 50% of attacks lasting less than nine minutes.

However, when we start to look at the extremes of our dataset and focus on the max values, we find a relatively large growth in capability. Compared to last year, the largest attacks increased by 198% in bits per second (BPS) and 156% in packets per second (PPS).

We don't normally spend too much time discussing the outliers because they are just that – outliers – but this growth in capacity from the attackers shows that they are continuing to build their abilities, even if the majority of their attacks tend to show similar characteristics as previous years.

Syncing with DDoS frequencies

One thing we wanted to understand a bit more this year was about the nature of the attacks. Not so much what and how, but the frequency in which organizations experience attacks and what that might tell us in terms of the nature of the adversaries who are targeting via DDoS. What we did was categorize organizations based on the time differences between attacks to help determine whether attacks were consistent, random or showed attributes of the attacks in bursts. What we found is that of the victim organizations, 2% had attacks that were periodic and relatively predictable, 40% of organizations had attacks that showed up as bursts of activity, and 57% of organizations experienced attacks that occurred in seemingly random intervals.

Figures 67 and 68 provide the breakdown of the differences and similarities that exist among our three classes of DDoS victims. For the victims that largely experience Bursty (aggressive and frequent) types of attacks, the median time between attacks was about one day. This is a pretty stark contrast to the Random (no discernable patterns) organizations that experienced about 14 days without an event.

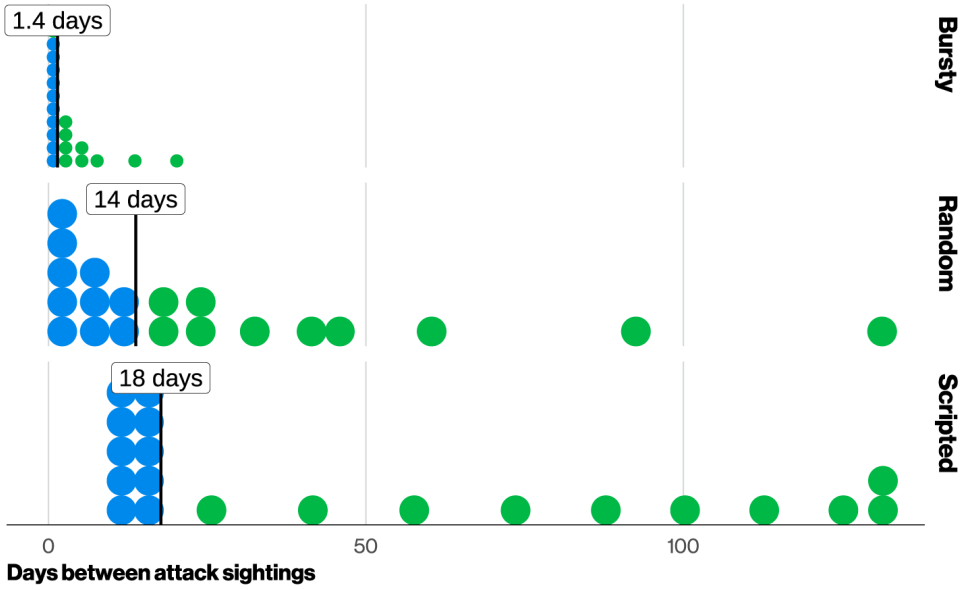

Figure 67. Distribution of time between attacks based on pattern (n=217)

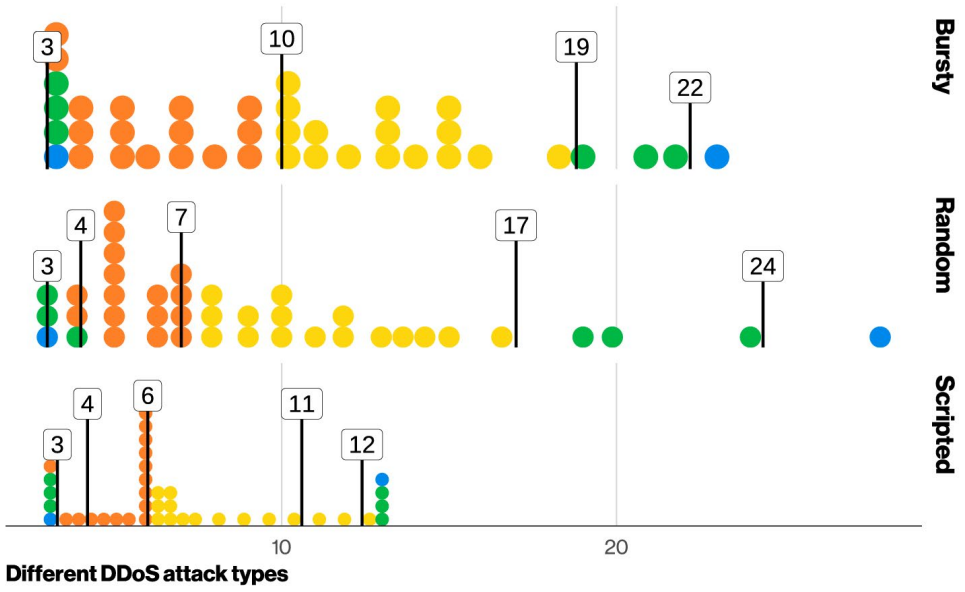

Figure 68. Distribution of number of different attack types (n=217)

In addition, over the course of the year, our Bursty victims would face about 17 different attacks (with a median duration of 24 minutes), while our Random victims had a median of six events in the same time period (also with a median duration of 24 minutes).[84] The higher end of the distribution for duration is around 88 minutes for the 90th percentile and 101 minutes for the 95th percentile. The question arises, "What would be the impact to your business if your business was affected?"

Lastly, when it comes down to the types of DDoS attacks experienced, organizations have to be able to mitigate a complicated concoction of different techniques. Attackers commonly target organizations with at least 10 different attack types tailored to different types of network protocols and services over the course of the year. For organizations that experienced more Bursty styles of attacks, 50% of victims experienced at least 10 different types of attacks, in comparison to the other attack type victims experiencing only seven.

As attackers continue to build their capabilities and capacity for their attacks, we as defenders have to maintain our footing and use collaboration and partnerships to counteract these actors.

84. Common durations aren't super surprising to us as these different types of organizations are probably being targeted by the same attack toolkits, just in different frequency.

The (AI) bots are back in town.

GenAI in general, and agentic AI specifically, can potentially bring automation but also disruption to markets and supply chains. The DBIR can confirm that – as far as bandwidth considerations and protection of intellectual property are concerned – the disruption is possibly already here. Consequently, cybersecurity teams are scrambling to automate handling the surge of AI scanning bots, and AI-driven traffic in general.

There are two main categories to be found today: the AI crawlers that have a similar function to the old search engine crawlers and gather data for training and fine-tuning of models and the AI fetchers designed to act on a direct request from users (or indirectly from AI agents) to retrieve information to support a task. If you are feeling too lazy to read a website and you ask the model to summarize it for you, the fetcher is there to gather the information.

The traffic generated by these bots has been growing significantly. Roughly 15% of non-malicious bot traffic in Q3 2025 was related to AI bots, whereas our good old search engine crawlers were responsible for 60% of the pie.[85] Think about that for a second. We have these new types of bots that did not exist four years ago, and that now represent as much as one-quarter of the traditional search engine crawler traffic.

And it keeps growing. With a bot traffic dataset spanning May 2025 to Dec 2025, we calculated the Compound Monthly Growth Rate (CMGR) of AI bot growth in general terms, but also by industry.

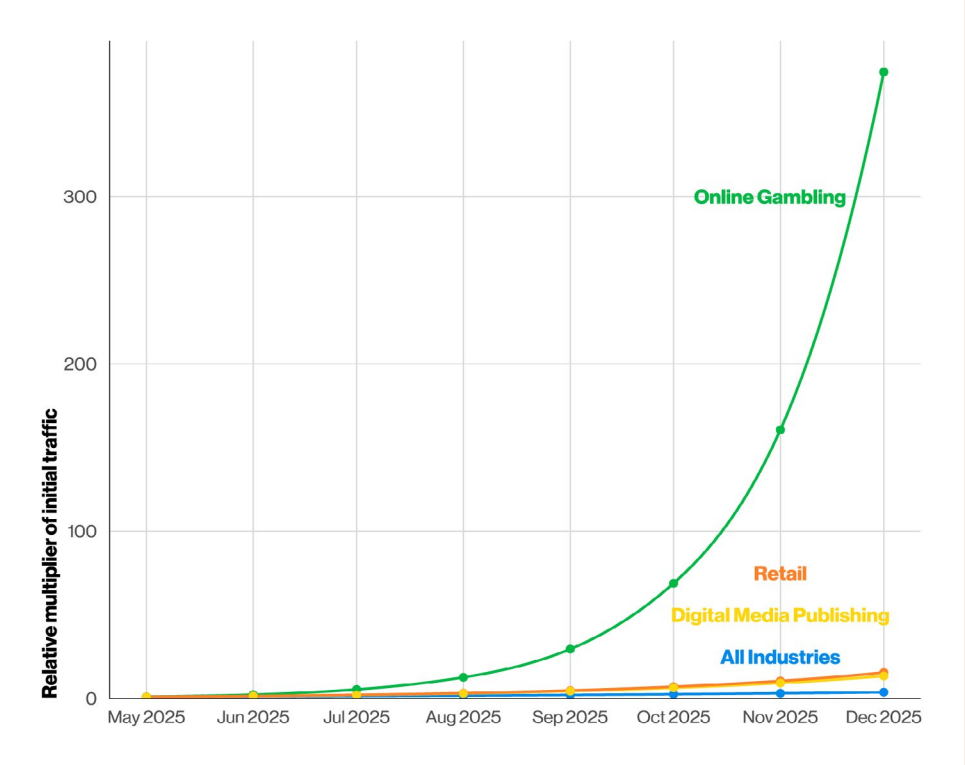

Figure 69. Relative growth of AI bot traffic over time

Figure 69 shows the surprising comparison of relative growth of all industries and the top three individually. During the reporting period, global traffic from AI crawlers and fetchers grew 21% month over month (MoM) across all observed industries, a 4% growth on fetchers and a 32% growth on crawlers. For reference, human-led traffic growth was essentially flat at 0.3% CMGR.[86]

The top three industries in growth tell us a lot about the objectives of these bots. Online Gambling has seen an absurd growth of 133% MoM,[87] demonstrating how quickly AI systems can vacuum up high-value, data-rich environments.

The next two, Digital Media Publishing and Retail organizations – with growth rates between 45% and 48% – have business models that depend on direct traffic and engagement, raising concerns regarding content authority, attribution and commercial impact.

Managing this increased traffic will be a challenge, doubly so if you care about the dichotomy between crawlers and fetchers. Accounting for increased resource usage in this evolving landscape is the bare minimum, and bot management solutions would be required for more fine-grained control, especially if your content is proprietary and monetizable.

85. Findings from our research partner Fastly. You can find the research referenced and their latest reports here: fastly.com/threat-insights.
86. This is a good time to brush up on the Dead Internet Theory: en.wikipedia.org/wiki/Dead_Internet_theory.
87. We are surprised the prediction markets didn't predict that.

Deep-dive analysis

/04

The paths of privilege escalation

Passwords, configurations, permissions and patches

Although the title has a relatively nice ring and sing-songy element to it,[88] it might also strike a familiar chord of "isn't this just what cybersecurity is?" And yes, we could probably understand cybersecurity as largely residing within these four pillars. What's important, though, is that this gives us a relatively simple communication vehicle and method of determining which ones we believe are important in which context. This can be how many successful ransomware attacks, how many red team engagements and how this contrasts to the daily hamster wheel of chasing down critical alerts and findings.

Over the years, we have developed an understanding of the different initial access vectors, and now the questions we have are: How are adversaries moving from low-level users to domain admins, and are we as a community prioritizing these gaps appropriately?

The baseline

The best starting point for this type of analysis is understanding the foundation. In this case, we will leverage the MITRE ATT&CK[89] model. For those who aren't experts on it, no worries; we aren't going to be splitting hairs between "Path Interception by Search Order Hijacking" and "Path Interception by PATH Environment Variable." All we are really focused on is that ATT&CK provides a lingua franca of what different types of actions (techniques) actors use to achieve their objectives (tactics).

Although there are more than 691 different techniques as of ATT&CK version 18, what we're really going to be focusing on are two of the Navigator columns: Privilege Escalation and Credential Access. There are (only) 176 techniques, as these are the ways in which low-level users are either able to escalate the permissions on or compromise higher-privileged accounts and are a key step in many attacks.

As part of each of these techniques, there are a set of mitigations,[90] and each of these mitigations are grouped into one of the four buckets we just discussed (passwords, configurations, permissions and patches).

Let's do a simple level-set of what we mean by each element:

1. **Passwords:** Pretty obvious as to what we mean by passwords, but what we want to highlight is the use of crackable or guessable passwords. In the world of Windows, hashes are heavily utilized as the means of authentication, and weak passwords in combination with weak hashing protocols are often a simple recipe for attackers getting a hold of the actual credentials.
2. **Configurations:** These are the toggles and settings that exist in operating systems and services. Our point of reference was to refer to the CIS Benchmarks[91] as our source of truth in terms of what are considered configuration elements.
3. **Permissions:** Permissions are the rights allocated to users, services and groups. A simple example of a permission is "SeDebugPrivilege" in Windows environments, which is typically assigned to the administration local group and allows administrators to attach a debugger to any process and is commonly leveraged to dump credentials from memory.
4. **Patches:** Of course, outdated and vulnerable systems and some vulnerabilities can give attackers a straight shot to credentials or higher permissions.

88. An early draft had "configurations" as "parameters" in the title, but that was trying a bit too hard.
89. attack.mitre.org
90. Except for ~10% of techniques that don't have mitigations listed for some reason. Ah, maybe in version 19.
91. A community resource produced by the non-profit Center for Internet Security, which defines hardening security guidance among other helpful community-focused resources

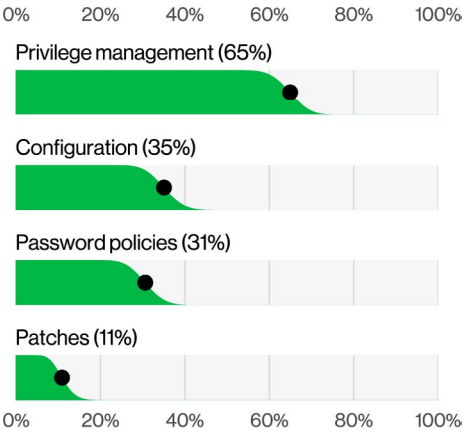

Figure 70. Percentage of techniques mitigated by mitigation type (n=157)

With that explained, let's take a quick look at which techniques are mitigated (partially or entirely) by each type of safeguard in Figure 70. As you can see, the biggest percentage of techniques can be addressed by Privilege Management, with about 65% of techniques showing this element as being key to mitigating the attack. The types of mitigations we find recommended include restricting administrative functionality or restricting access to sensitive folders, such as startup folders or top-level operating system directories.

The next two most common mitigations are Configurations and Password Policies, at 33% and 30%, respectively. These types of mitigations focus on tweaks to the operating system and servers with the recommendations pertaining to password practices, such as password strength, MFA and disabling insecure protocols or weak hashing algorithms.

Lastly, we see that only about 10% of our techniques are mitigated by applying patches. It's worth pointing out that Exploitation of vulnerabilities is far more traditionally present in incidents as an initial access vector (as seen in the "Results and analysis" section of this report) and will only have limited presence here in Privilege Escalation. And though this starting point is still a relatively large field of things that can be done, we will narrow our analysis to the techniques that are being leveraged.

Understanding observations

The first place we want to look is how we can quantify the commonness of each technique. To get an understanding of prevalence, we combined a variety of different sources, each with their own set of biases and limitations, and found out how this wide view can help us prioritize our security efforts.

- **Threat intelligence:** One way that we can understand observations is by examining threat intelligence reports, which are recorded observations of different techniques being used by Actors.
- **Red team activity**: Red teams are ethical hackers that test an organization's ability to detect and respond to attacks. Some of these tests are structured as "Assume compromise," in which attackers are given access to a system within the environment and leverage that system to conduct their attacks, rather than having to establish a foothold on a compromised system.
- **Incidents:** The incidents that we're examining here are a slightly different subset from our larger incident corpus. Some of these incidents are incorporated into our larger dataset and some are ineligible for inclusion. The reason we're examining this subset of incidents is that a few of our data contributors provided us with a sufficient level of detail to allow us to accurately map things to ATT&CK and VERIS.[92]

Figure 71 captures the techniques found across these different data sources. An interesting finding is OS Credential Dumping[93] (T1003), specifically LSASS dumping, was one of the most common techniques found in both the threat intelligence dataset (34%) and the incidents dataset (20%), as shown in the figure.

The limited appearance of this technique in red team engagements is probably more due to the structure of the testing ("Assume compromise") rather than it being a technique beneath the hack-o-mancers that perform these tests. LSASS memory is one of those interesting cybersecurity problems. It's foundational to how Windows works, and there are built-in mitigations in place; however, it's still the bread and butter[94] of attackers once they get access to an endpoint system.

In contrast, in the red teaming world, we see a lot of targeting of authentication services by using Kerberoasting[95] attacks and stealing and forging authentication certificates. These types of attacks are largely focused on compromising service accounts that are misconfigured and have weak passwords.

92. To be clear, most of our data contributors do so, and we love them. However, some of our data comes from public sources, where this level of detail is harder to come by.
93. Semantics alert: Because we're incorporating a bunch of different datasets, we aren't capturing or representing hierarchical relationships here that exist between the techniques and subtechniques.
94. Or the "meat and potatoes" or the "rice and beans," depending where in the world you are from
95. Apologies for the ultra-oversimplification of these attacks, but we got to stick to our report page budget!

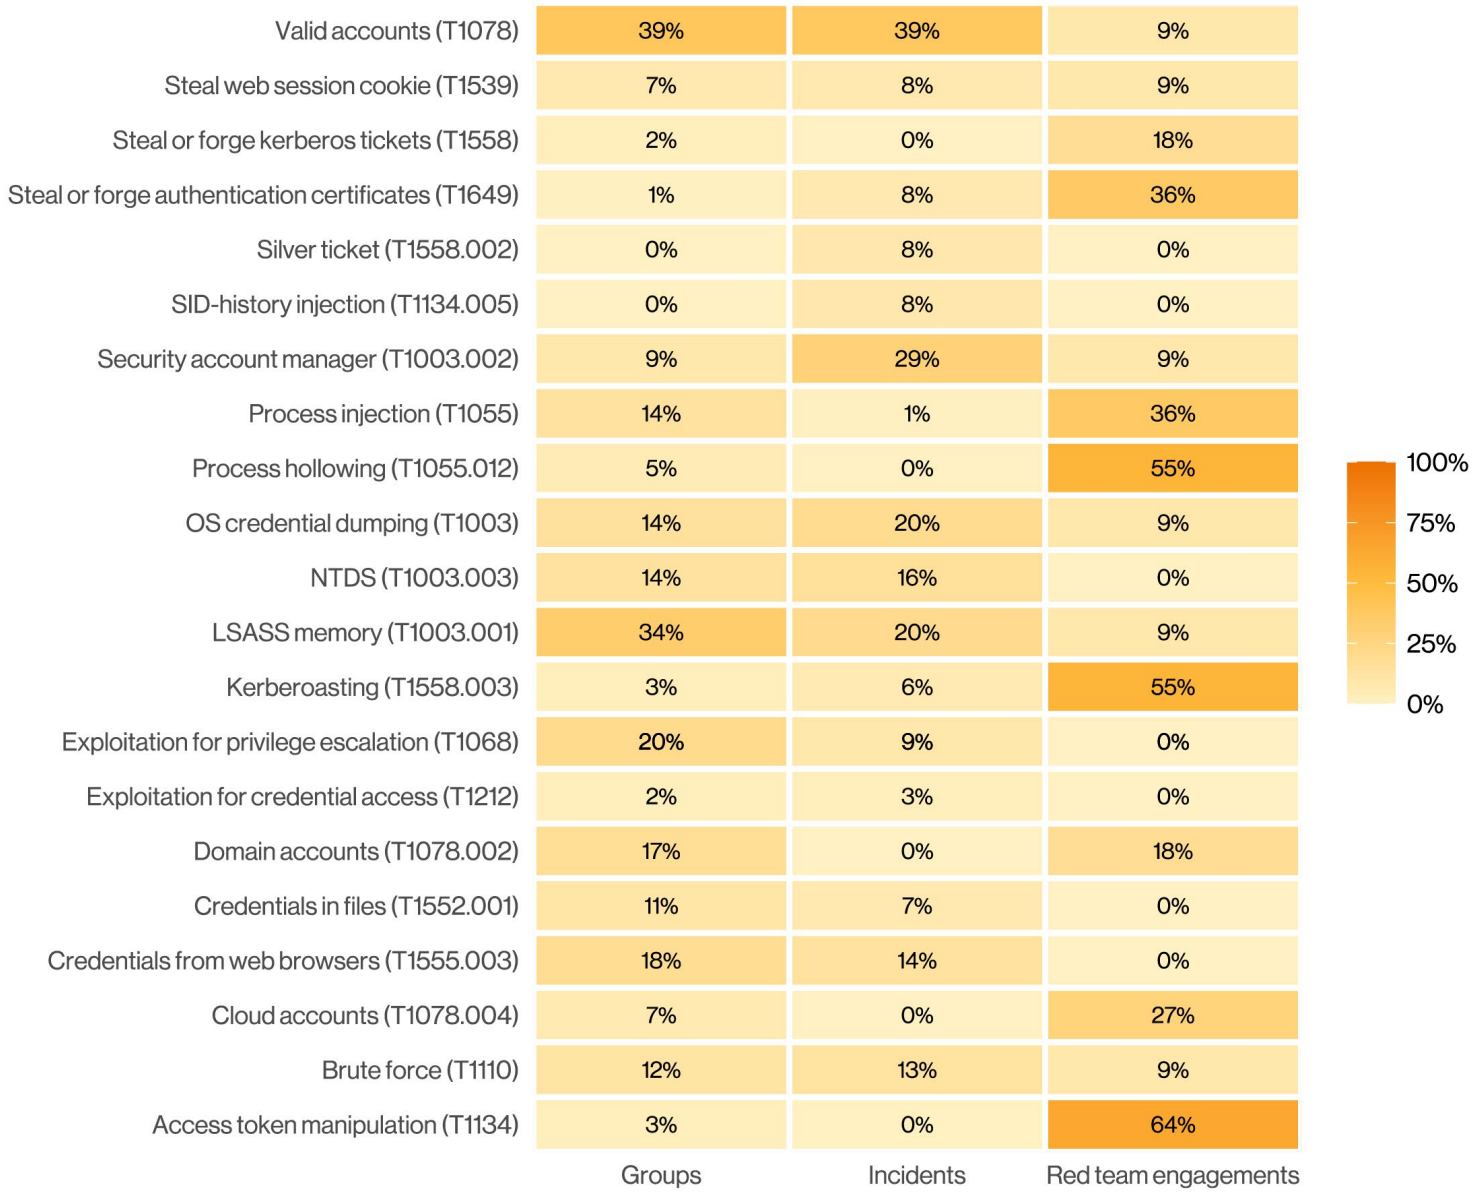

Figure 71. Privilege and credential access techniques by observations

We also see the attack DCSync and Security Account Manager (SAM) dumps showing up, which are alternative methods of dumping credentials. DCSync does this by convincing a remote domain controller to replicate its data (including passwords), and SAM dumps extract credentials either from memory or from a secure location in the registry.

Lastly, it's worth mentioning Exploitation for Privilege Escalation, which shows up in 9% of incidents and 20% of threat intelligence. That means the majority (83%) of incidents did not include exploiting a vulnerability[96] for escalating privileges. So even though patching will have a positive impact on your initial access vector mitigation, it's not the only way to move the needle for preventing privilege escalation. However, when we took a quick look across our vulnerability management dataset for this year, we found 26% of organizations still had privilege escalation vulnerabilities from 2021, and 11% had some from 2018. Let's focus on our other mitigation strategies, but don't leave those old vulnerabilities behind.

Passwords

Is this the year passwords die? We say this in jest, nearly every year, knowing full well in our heart of hearts that's likely not the case. Passwords are used in various places as the only line of defense. However, this line of defense may be more of a Maginot Line than The Wall of Westeros.[97] When it comes to passwords, we as an industry have previously tried to drive complexity recommendations as the foundation of making strong and uncrackable passwords.

However, when it comes to password strength, the battle isn't just against compute power (at least not until the quantum singularity[98] occurs) but also against human nature.

Remembering and creating unique passwords is challenging, but enforcing complexity isn't. We found that the median percentage of Active Directory user accounts not meeting complexity requirements was less than 1%.

Although organizations are good at enforcing "strong passwords" – it's a default option in most systems nowadays – that is not keeping users from using passwords that have already been compromised, with the median percentage being 4%, as Figure 72 demonstrates.

To put it simply, based on this dataset, users are more than four times more likely to use an already compromised password than a "weak" password.

Compounding this issue, the median percentage of users that are reusing passwords, or have the same password as others, is about 6%. When attackers are looking to crack the hashes of the passwords they have collected from their engagements, they are banking on the fact that passwords are being reused, either from previous compromises in other organizations or internally. If the systems also aren't configured to use modern hashing algorithms, it makes their job even easier, which is a good segway to our next section.

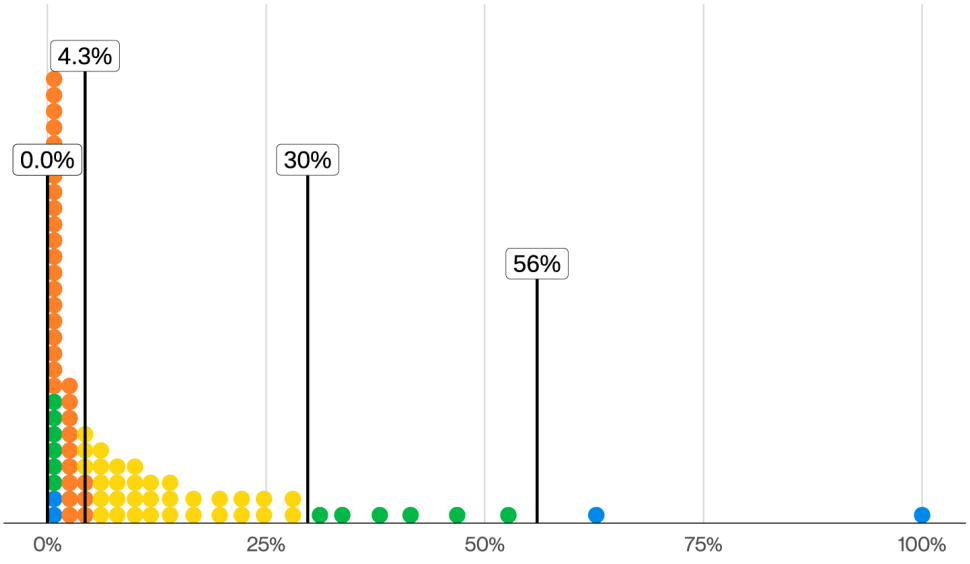

Figure 72. Distribution of percentage of accounts scanned with breached passwords (n=7,345)

96. Vulnerability in the strictest definition as being a CVE-assigned vulnerability. Yes, that does include the ones in the backlog of the National Vulnerability Database.
97. What do you mean The Wall fell in Season 7? Game of Thrones only ran for six seasons. That's not funny.
98. With or without AI. Take your pick.

Configurations

We are fully aware that security configurations aren't the spicy and exciting part of cybersecurity.[99] Turning on features, testing them in a staging environment, pushing them to all employees, reverting them back to the previous state after some ancient application breaks, crying in the office bathroom and calling it a day doesn't really make for the most compelling storytelling.

However, secure configurations are a key component to preventing many of the techniques that we have discussed in this section. Of those, about 33% have some element of configuration being listed as a key mitigation.

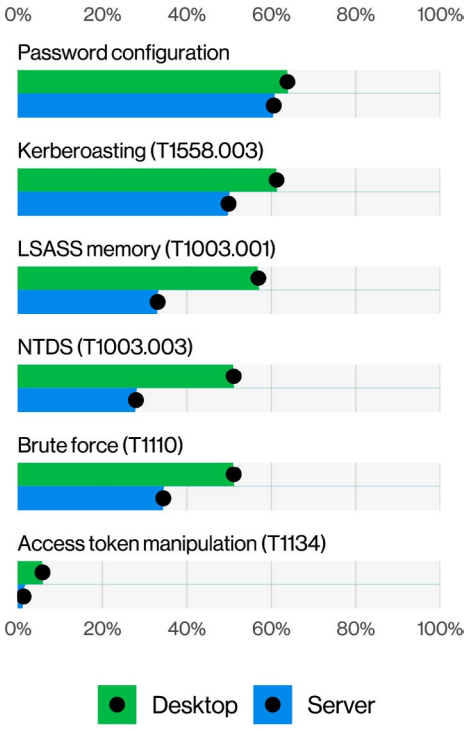

Figure 73. Percentage of failed configuration checks by attack techniques and asset type
(n for Desktop=10,279,021)
(n for Server=1,149,990)

The solutions exist, but they often require the organizations to explicitly implement those, as vendors have to balance the business need of making their software easy to use out of the box and these additional security risks.

Figure 73 has a breakdown of the percentage of failed configuration checks associated with different ATT&CK techniques as compared between servers and desktops. First, some major takeaways: The majority of assets failed password configuration checks based on traditional security hardening recommendations. The most common failed checks related to passwords are the number of failed login attempts before lockout (97% of assessed devices) and passwords requiring lengths of 15 characters or more (90% of assessed devices).

The most implemented configuration, in which less than 10% of assets failed the check, pertained to protecting against Access Token Manipulation. These types of techniques rely on systems that are misconfigured and allow users with access to administrative functions. Also of note is the difference in LSASS protection between servers and desktops. Desktops continue to be largely vulnerable to these types of attacks in comparison to servers, which when we consider the possible exposures that exist from administrators logging into end-user systems, constitutes a perfect segue for our next section on permissions.

Permissions

The longer you spend staring into ATT&CK techniques,[100] the more you start to see a common pattern: the requirement for proper permissions to conduct the attack. Having access to permissions of a local admin (or, even better, domain admin) opens up several venues for attackers, such as reading into protected files and memory or making key configuration changes to a system. Of the techniques we examined, 65% mentioned restricting permissions as one of the main ways to prevent these techniques.

Privileges are a bit more nuanced than one might think, though, and don't simply involve looking at a list of objects that have overly strong permissions. In some instances, these permissions are transitory in nature. One privilege may allow you access and/or allow you to modify an object that might, in turn, allow you more privileges. These are sometimes understood as attack paths, whereby chaining these transitory relationships of privileges, attackers are able to pivot from low-level users to domain admins. To understand the complex relationships that exist within an environment, organizations can map the relationships that exist between users, groups and permissions to create an attack graph. These attack graphs can show how attackers find the easiest paths between low-level users and domain admins and assist defenders in identifying exposures. Based on attack graphs collected from organizations, it was found that 16% of organizations had about 80% exposure, meaning that given initial access to the environment, an attacker with low-level privileges had an 80% or better chance of successfully compromising a key administrative account or infrastructure element.

99. Is there one?
100. The techniques eventually stare back.

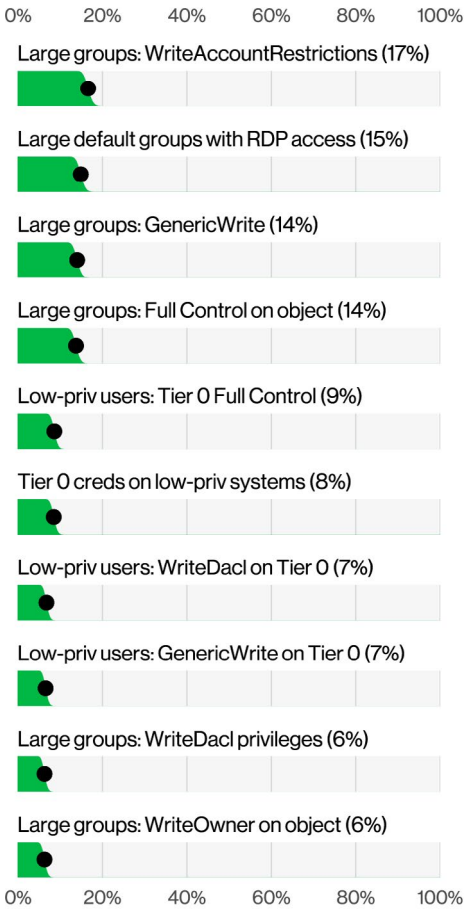

Figure 74. Top attack paths by permissions (n=1,500)

Figure 74 captures the 10 most common attack paths found in these assessments. Some of these exposures, such as Tier Zero accounts logging into non-privileged systems, could directly expose high-privilege accounts to techniques such as Credential Dumping. Still, other privileges can open up organizations to different types of attacks, depending on how the systems are configured and the type of access the attackers have, such as the ability of low-level users to have write privileges over higher-tier resources, a clear compromise of the least-privilege principle. For organizations to help properly defend themselves, they need to adopt a stronger stance than just chasing down vulnerabilities and alerts. Instead, they need to examine their cybersecurity in a holistic fashion, including examining the exposures that exist in their environments from permissions, configurations and poor passwords.

The North Korean IT worker risk

Overemployed and over-trusted

One of the defining features of the 2025 IT landscape was the systematic infiltration of IT workers (ITWs) from the Democratic People's Republic of Korea (DPRK). Some of these workers – which we shall refer to collectively as ITWs – were able to achieve multiple positions across a large spread of industries.

Using stolen identities, the ITWs were able to acquire jobs and operate out of regionally hosted laptop farms run by local accomplices. This setup allowed the actors to pass the interview process and perform the jobs without requiring a physical presence in the area. Interestingly enough, some of these ITWs did more than just "quiet quit" and collect a paycheck, as some organizations were surprised to find that some of their top-performing new recruits were misrepresenting their identities.[101]

This is a pretty serious concern on the outset, since employing individuals with falsified credentials and skills poses a big enough risk, but adding the fact that these employees may also be associated with or operating on behalf of the North Korean government, according to U.S. government advisories and law enforcement filings, is enough to give your compliance and security officers a collective heart attack.

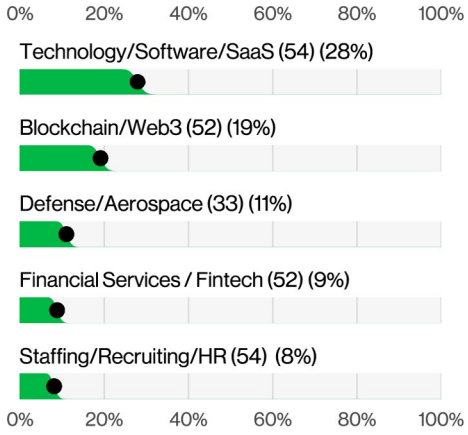

Figure 75. Top targeted industries by North Korean ITW campaigns

I don't have anything interesting. Why would I be a target?

When tackling a complex issue, it's helpful to get an idea of the scope of the problem. In 2025 and earlier, there were various reports[102] and indictments[103] that came out from law enforcement as they continued to actively pursue and disrupt these operations. Reviewing those sources, it was clear that companies from widely different sectors were falling victim to these types of fraud.

Figure 75 shows the variety of different industries that were being targeted in 2025 by ITWs. It seems that the vast majority of the activity uncovered was focused on the raising of funds through payroll. Some industry sources have suggested that ITWs may facilitate further access for state-sponsored groups, though we have not been able to substantiate this with sufficient data.

Competitive job market challenges

For obvious reasons, ITWs have historically targeted remote jobs focused on programming and data engineering. Given they have to pass the technical interviews and make a convincing enough appeal to be hired, job specifications needing IT and software development skills have been the primary focus.

As of more recent times, with shifts in the hiring market and more awareness of the specific jobs that would be targeted, they have also started to move toward human resources and marketing jobs.

Figure 76 has the breakdown of job specifications that have been targeted by ITWs in 2025 and include frontend developers, with blockchain/Web3 and full-stack engineering jobs at the top. Other researchers[104] that track this threat have found these actors pivoting to AI-focused jobs, as well, showing that they are following market trends in terms of where those high-paying remote jobs are. Too bad those folks are not active on LinkedIn, as we could all benefit from some hiring tips in hot markets.

101. This sounds like a traditional DBIR joke, but this fact was raised to us in multiple interviews with subject matter experts (SMEs) we conducted while writing this section.
102. reports.dtexsystems.com/DTEX-Exposing+DPRK+Cyber+Syndicate+and+Hidden+IT+Workforce.pdf
103. justice.gov/opa/pr/justice-department-announces-coordinated-nationwide-actions-combat-north-korean-remote
104. okta.com/blog/threat-intelligence/north-korea-s-it-workers-expand-beyond-us-big-tech

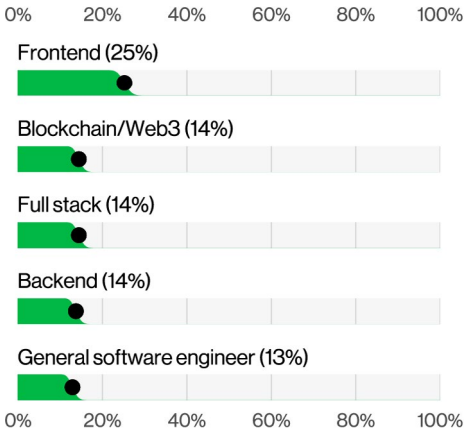

Figure 76. Top targeted job specifications by North Korean ITW campaigns

How big of a problem is this?

Reaching a precise number of victims is relatively challenging, as organizations are not required to publicly disclose that they have accidentally hired a North Korean IT worker, but looking at the other side of the equation can give us an understanding of how many fake ITWs there are. According to our analysis, ITWs leveraged an estimated 15,000 possible stolen identities, with the typical ITW leveraging around three to five identities at any given time. With those back-of-the-napkin calculations, our rough analysis suggests the figure could be in the low thousands, though this estimate carries considerable uncertainty.

In terms of protecting your organization, like so many other things in security, it comes down to collaboration within your organization, especially with the teams running human resources and recruitment processes. That includes things like:

- Additional scrutiny to backgrounds, resumes and information provided by applicants
- Verifying identity through multiple touchpoints during the hiring process
- Making sure your insider threat and security awareness programs discuss these new types of threats

Industries

/05

Introduction

Welcome to the "Industries" section of the 2026 DBIR! As noted previously, we analyzed more than 22,000 confirmed data breaches this year – by far the largest number we have ever examined in a single report. In this section, we break those breaches down by industry. Different industries face different threats, largely because their attack surfaces are not created equal.

When reading this section, it's important to be aware of a few caveats. Industry-level differences can be influenced by factors such as varying regulatory and reporting requirements and the resulting differences in the level of external scrutiny they receive, along with the size of the data sample we have for any given industry. These and other factors can affect how a vertical appears in the report, so please keep that in mind when comparing one industry to another.

Many of our readers refer to this section to find "bespoke post" results for their own industries. When doing so, we recommend starting with the top patterns for your sector in this section, then circling back to the main pattern chapters referenced for your vertical. This should give you a deeper understanding of the attacks you will likely need to defend against.

Finally, since this is a security report and not a novel penned by Tolstoy or Proust, we do not have sufficient space (and in some cases not enough data) to look at all industries exhaustively; therefore, we provide Table 3. It offers a quick reference for high-level information on the industries that we do not delve into here.

	Incidents				Breaches			
Industry	Total	Small (1–1,000)	Large (1,000+)	Unknown	Total	Small (1–1,000)	Large (1,000+)	Unknown
Total	31,861	7,257	528	24,076	22,625	7,153	466	15,006
Accommodation (72)	319	89	11	219	250	88	11	151
Administrative (56)	422	295	12	115	419	295	12	112
Agriculture (11)	223	34	0	189	219	34	0	185
Construction (23)	843	525	8	310	828	524	8	296
Education (61)	1,302	219	22	1,061	1,252	216	22	1,014
Entertainment (71)	587	90	5	492	483	90	5	388
Finance (52)	3,809	365	52	3,392	1,300	358	50	892
Healthcare (62)	1,492	472	21	999	1,438	466	20	952
Information (51)	1,703	214	79	1,410	1,099	202	59	838
Management (55)	103	30	0	73	101	30	0	71
Manufacturing (31–33)	3,627	1,198	56	2,373	2,713	1,176	50	1,487
Mining (21)	72	35	3	34	70	35	2	33
Other Services (81)	900	206	4	690	885	205	4	676
Professional (54)	3,578	1,400	82	2,096	2,558	1,380	58	1,120
Public Administration (92)	3,634	148	17	3,469	2,410	148	17	2,245
Real Estate (53)	505	189	1	315	499	188	1	310
Retail (44–45)	997	315	32	650	806	313	32	461
Transportation (48–49)	689	242	24	423	652	242	21	389
Utilities (22)	638	57	8	573	597	57	7	533
Wholesale (42)	1,057	877	18	162	1,048	875	18	155
Unknown	5,361	257	73	5,031	2,998	231	69	2,698
Total	31,861	7,257	528	24,076	22,625	7,153	466	15,006

Table 3. Number of security incidents by victim industry and organization size

Incidents

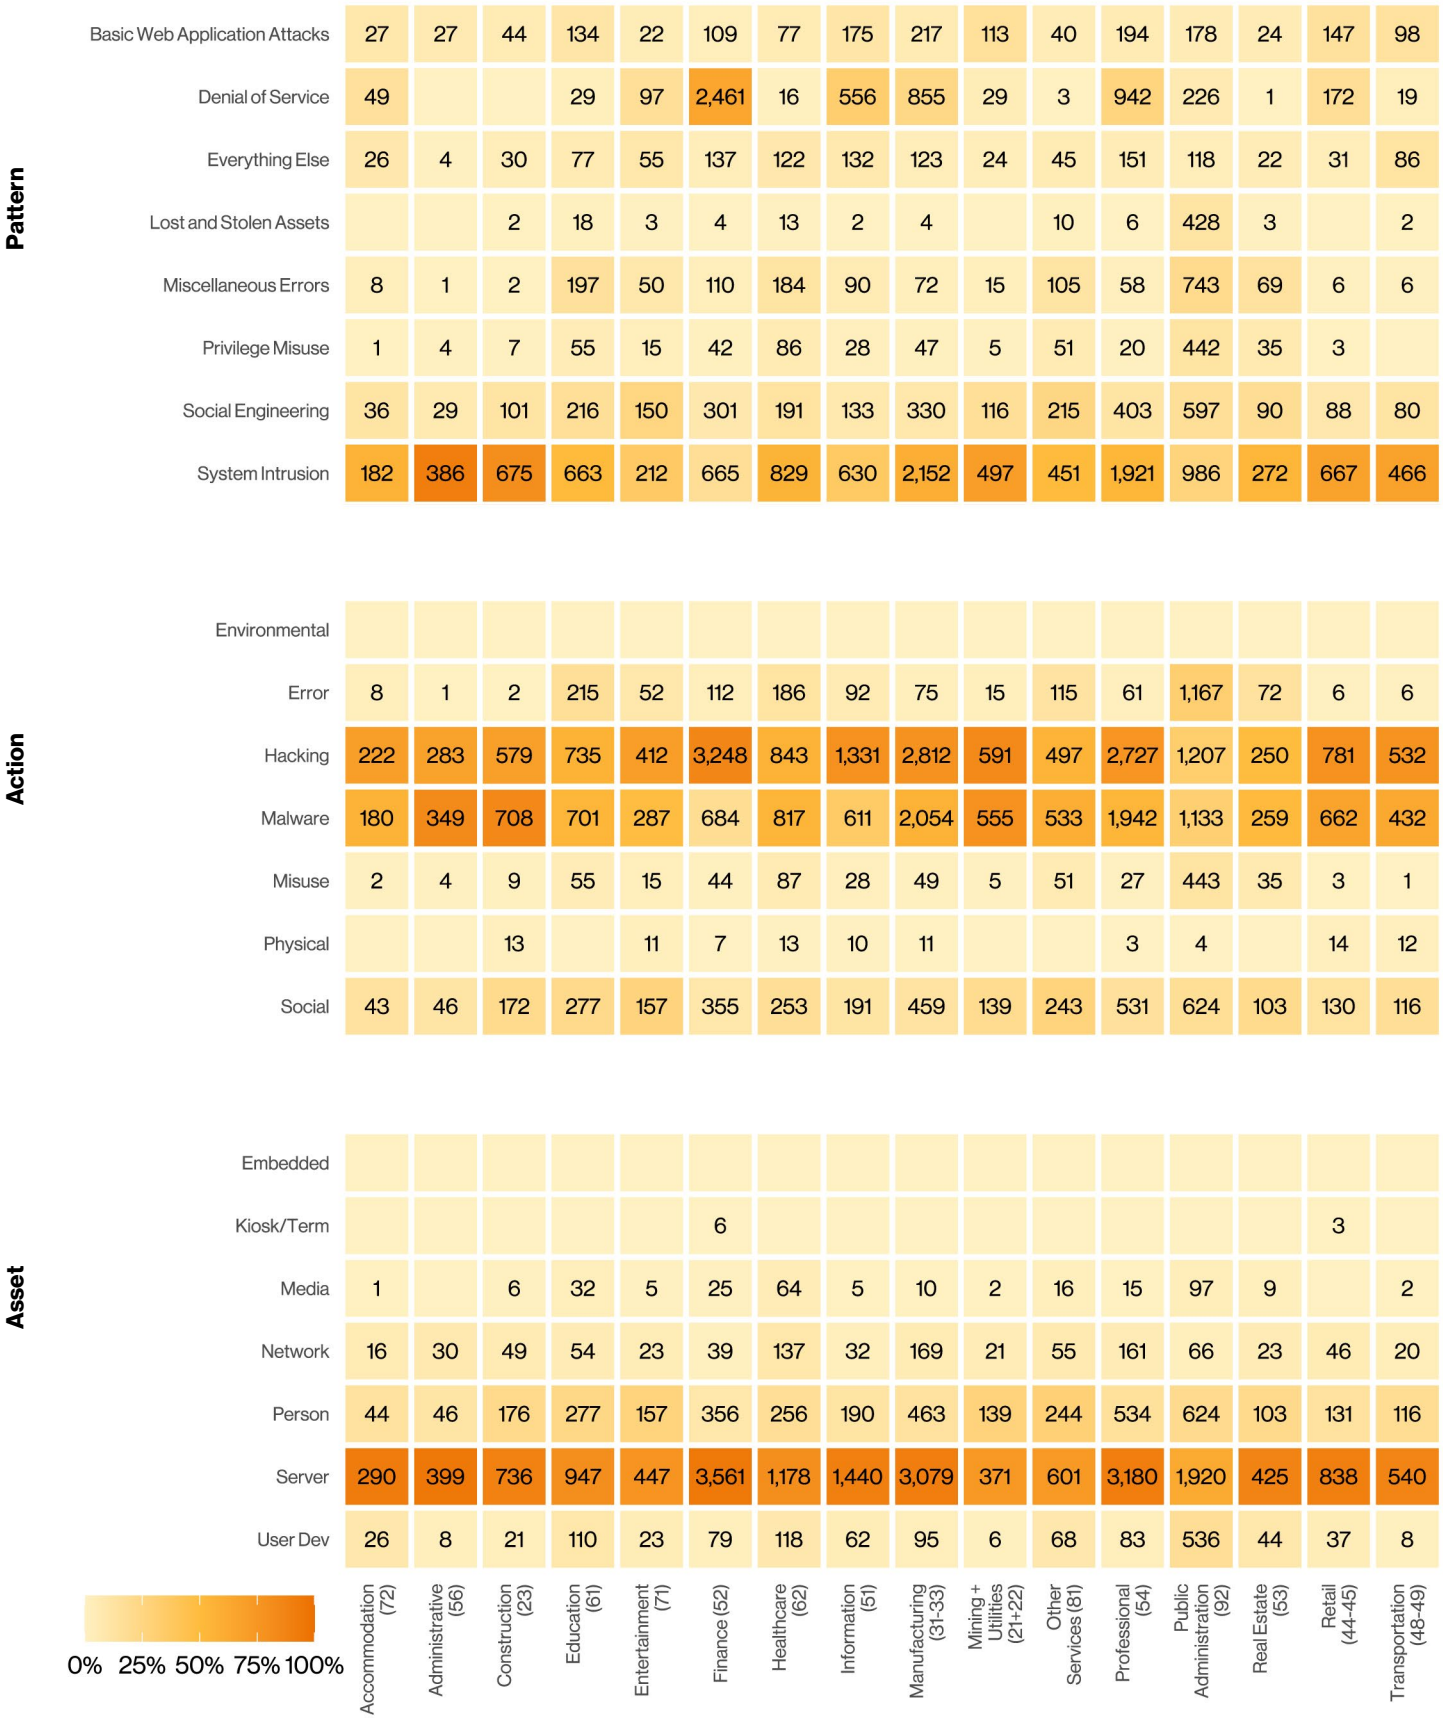

Figure 77. Incidents by industry

Breaches

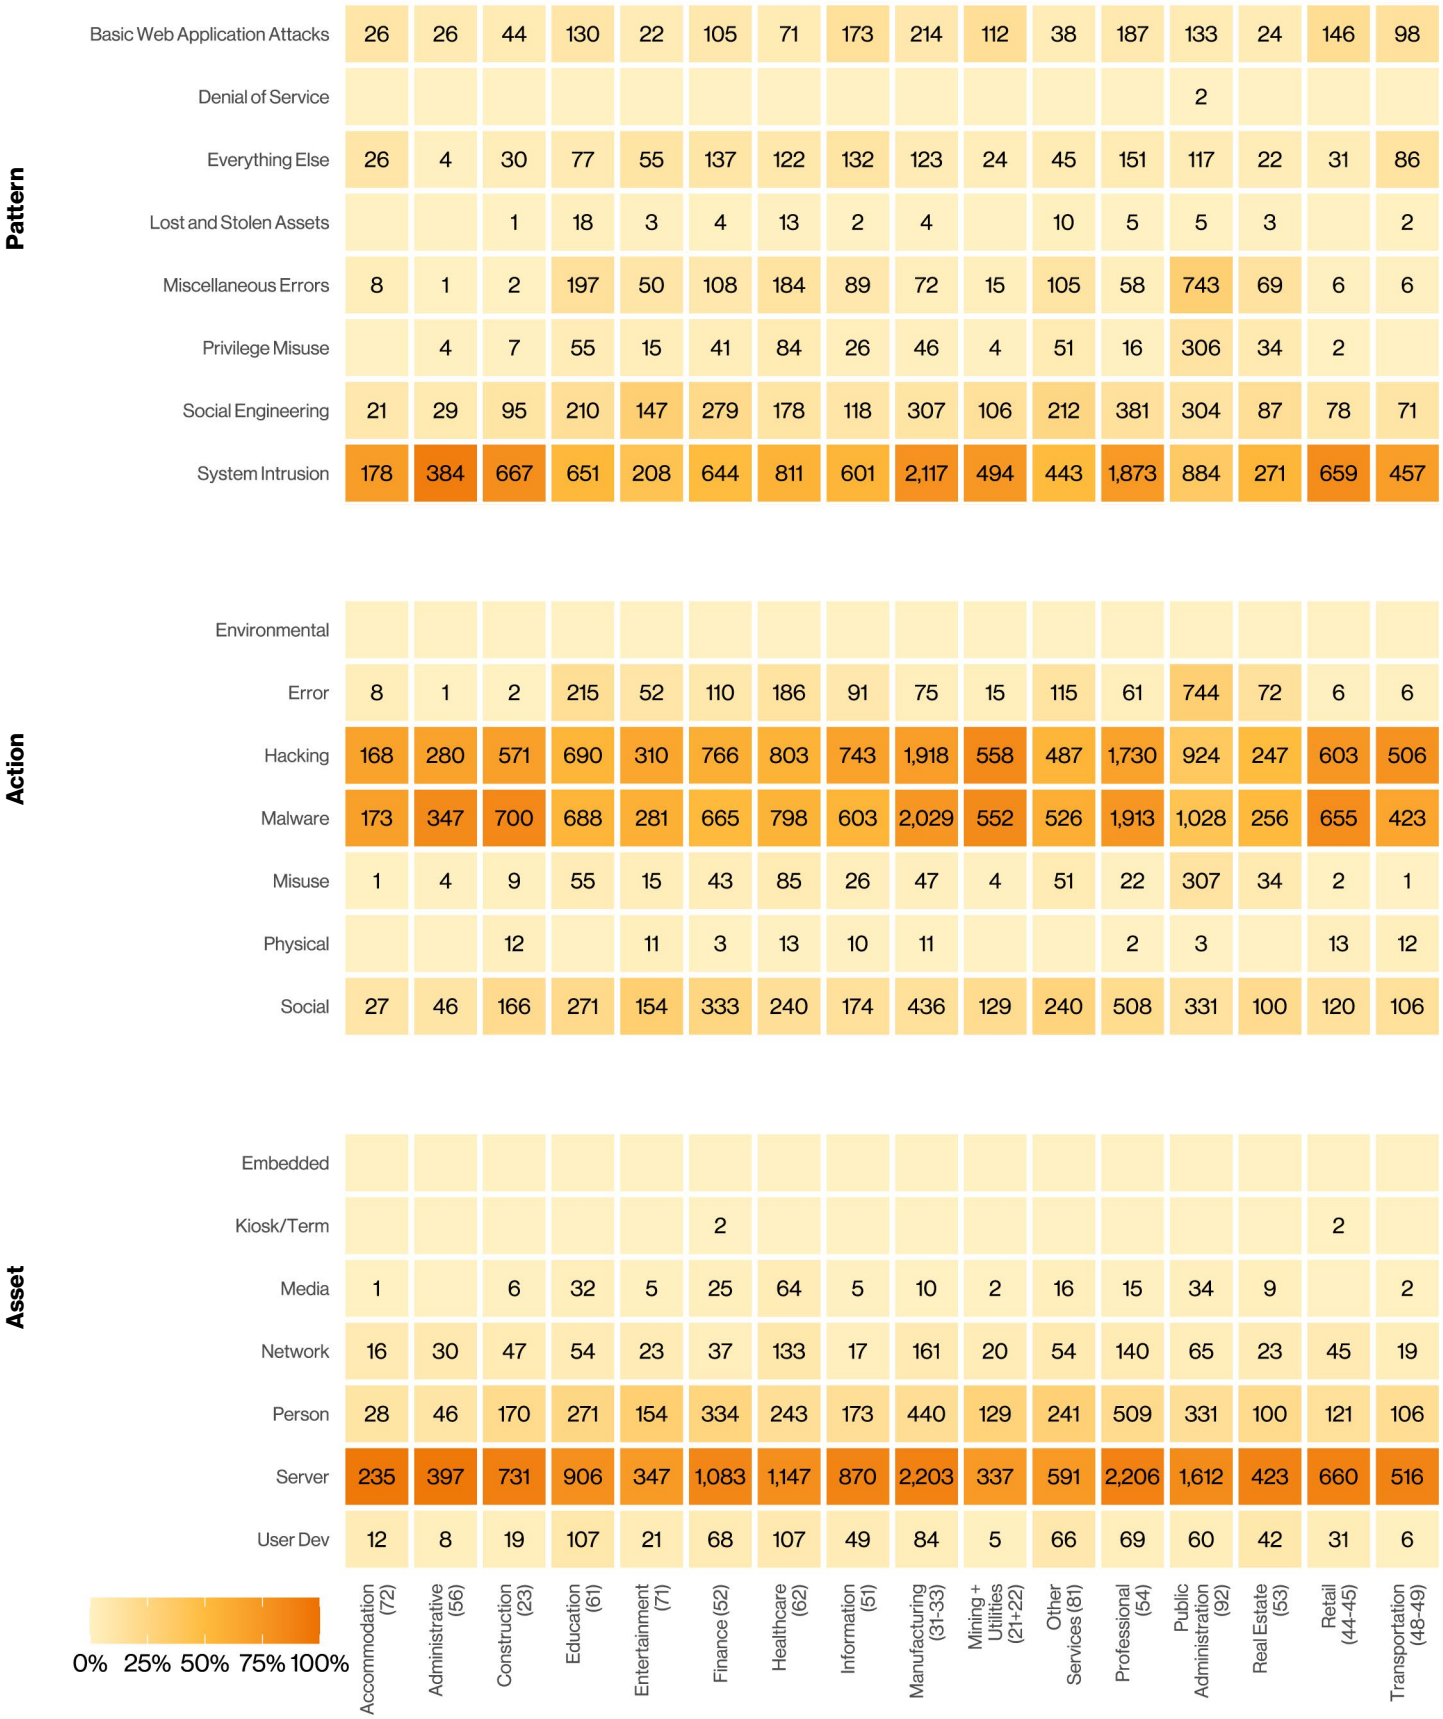

Figure 78. Breaches by industry

Industry (NAICS)	Frequency	Top patterns	Threat actors	Actor motives	Data compromised
Agriculture (11)	223 incidents, 219 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Social Engineering represent 91% of breaches	External (100%) (breaches)	Financial (71%), Espionage (29%), Ideology (1%) (breaches)	Internal (70%), Other (43%), Secrets (36%) (breaches)
Administrative (56)	422 incidents, 419 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 98% of breaches	External (99%), Internal (1%) (breaches)	Financial (100%) (breaches)	Internal (96%), Credentials (28%), Other (2%), System (2%) (breaches)
Construction (23)	843 incidents, 828 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 95% of breaches	External (99%), Internal (1%) (breaches)	Financial (97%), Espionage (5%) (breaches)	Internal (86%), Credentials (34%), Other (13%), Secrets (6%) (breaches)
Entertainment (71)	587 incidents, 483 with confirmed data disclosure	System Intrusion, Social Engineering and Everything Else represent 82% of breaches	External (86%), Internal (14%) (breaches)	Financial (89%), Espionage (20%), Ideology (1%) (breaches)	Internal (54%), Personal (45%), Other (31%), Secrets (20%) (breaches)
Information (51)	1,703 incidents, 1,099 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Everything Else represent 79% of breaches	External (89%), Internal (11%), Multiple (1%) (breaches)	Financial (84%), Espionage (16%), Ideology (2%) (breaches)	Internal (52%), Personal (39%), Other (31%), Credentials (24%) (breaches)
Management (55)	103 incidents, 101 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 98% of breaches	External (100%) (breaches)	Financial (100%) (breaches)	Internal (96%), Credentials (35%), Multi-factor credential (3%), Other (2%) (breaches)
Mining (21)	72 incidents, 70 with confirmed data disclosure	System Intrusion, Everything Else and Basic Web Application Attacks represent 96% of breaches	External (100%) (breaches)	Financial (97%), Espionage (1%), Ideology (1%) (breaches)	Internal (74%), Credentials (31%), Personal (17%), Other (9%) (breaches)

Table 4. At-a-glance table for victim industries without a section

Industry (NAICS)	Frequency	Top patterns	Threat actors	Actor motives	Data compromised
Other Services (81)	900 incidents, 885 with confirmed data disclosure	System Intrusion, Social Engineering and Miscellaneous Errors represent 85% of breaches	External (81%), Internal (19%) (breaches)	Financial (78%), Espionage (23%) (breaches)	Internal (66%), Personal (38%), Other (28%), Secrets (20%) (breaches)
Professional (54)	3,578 incidents, 2,558 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 91% of breaches	External (97%), Internal (3%) (breaches)	Financial (96%), Espionage (5%) (breaches)	Internal (80%), Credentials (31%), Personal (14%), Other (11%) (breaches)
Real Estate (53)	505 incidents, 499 with confirmed data disclosure	System Intrusion, Social Engineering and Miscellaneous Errors represent 85% of breaches	External (79%), Internal (22%) (breaches)	Financial (100%) (breaches)	Internal (63%), Personal (43%), Credentials (24%), Other (16%) (breaches)
Transportation (48–49)	689 incidents, 652 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Everything Else represent 89% of breaches	External (99%), Internal (1%) (breaches)	Financial (89%), Espionage (15%), Ideology (1%) (breaches)	Internal (84%), Credentials (27%), Secrets (16%), Other (14%) (breaches)
Utilities (22)	638 incidents, 597 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Social Engineering represent 94% of breaches	External (97%), Internal (3%) (breaches)	Espionage (71%), Financial (36%) (breaches)	Internal (85%), Secrets (68%), Other (21%) (breaches)
Wholesale Trade (42)	1,057 incidents, 1,048 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Social Engineering represent 99% of breaches	External (100%) (breaches)	Financial (100%) (breaches)	Internal (98%), Credentials (29%) (breaches)

Table 4. At-a-glance table for victim industries without a section (continued)

Educational Services NAICS 61

Summary

The Education vertical is troubled primarily by external, financially motivated actors who utilize Ransomware, Exploit vulnerabilities and rely greatly on the Use of stolen credentials.

What is the same?

The System Intrusion, Social Engineering and Miscellaneous Errors patterns are still the top three patterns, as they were last year and the year before.

Frequency	1,302 incidents, 1,252 with confirmed data disclosure
Top patterns	System Intrusion, Social Engineering and Miscellaneous Errors represent 83% of breaches
Threat actors	External (78%), Internal (22%) (breaches)
Actor motives	Financial (78%), Espionage (21%) Ideology (2%) (breaches)
Data compromised	Internal (64%), Personal (41%), Other (26%), Secrets (19%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (34%), Phishing (22%), Credential abuse (8%) (breaches)
Other metrics	Human element (68%), Third-party (40%)

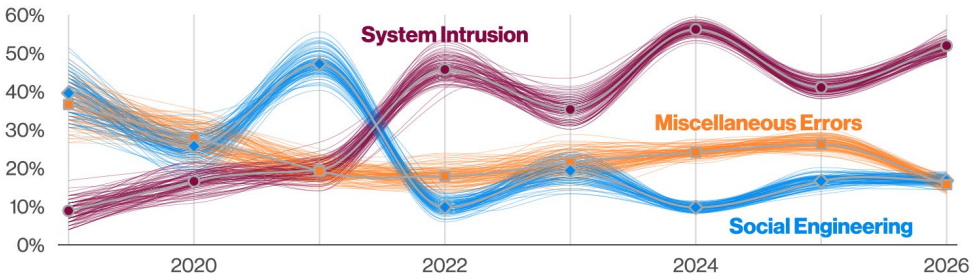

Figure 79. Top patterns in Educational Services breaches over time (n for 2026 dataset=1,252)

System Intrusion is the undisputed headliner for this vertical, appearing approximately three times as often as any other pattern and accounting for 52% of all Education breaches (Figure 79). In practice, that means many incidents involve attackers actively working their way into the victim environments – often by chaining together various action types or by generally employing whatever tools are necessary to pull off a more complex attack. Victims can take some consolation in the fact that they made the attackers earn every inch of the systems they compromised.[105] Small comfort, we know.

Social Engineering (17%) and Miscellaneous Errors (16%) represent significantly smaller numbers but still play meaningful roles. Think phishing emails that open the door or misconfigurations that kindly leave it propped wide open for anyone to walk through. Nevertheless, in this vertical, they are both only supporting roles and not the stars of the show.

The Education sector is beset by Hacking and Malware attacks in equal measure (55% of breaches), as shown in Figure 80. When we break down the Hacking varieties in play (see Figure 81), we get a slightly atypical picture. The Exploitation of vulnerabilities leads the pack, showing up in 77% of hacking-related breaches. Those weaknesses are being worked primarily by Organized crime and State-affiliated actors, although the data doesn't always go deep enough for us to name and shame the specific vulnerabilities that caused most of the damage. In a more familiar pattern, stolen credentials are also heavily represented, appearing in 65% of breaches that involve Hacking actions.

105. In true Chuck Norris style!

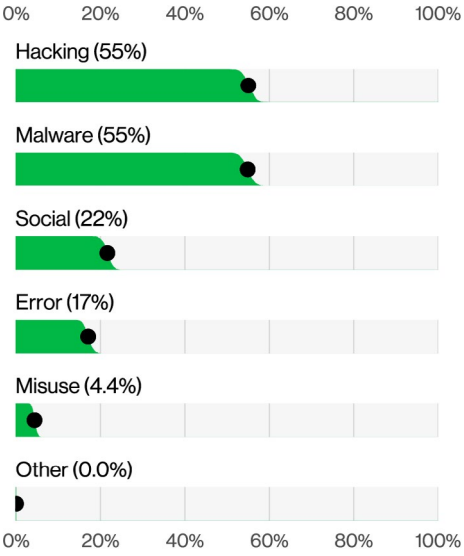

Figure 80. Action types in Educational Services breaches (n=1,252)

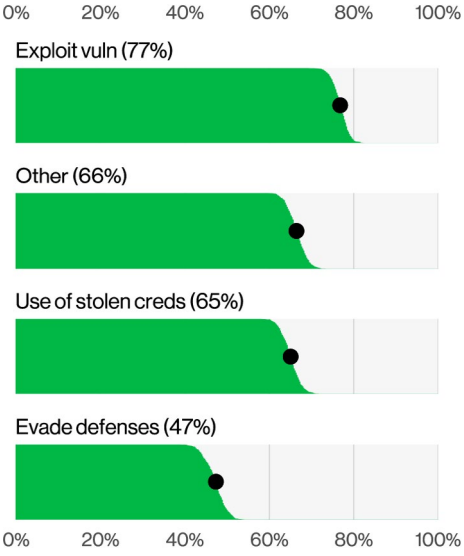

Figure 81. Top Hacking action varieties in Educational Services breaches (n=435)

Figure 82 confirms what most defenders in Education already suspect:[106] Ransomware remains the undisputed heavyweight of the sector, appearing in 65% of malware-related breaches. While this lack of a "plot twist" may not surprise seasoned veterans, the mechanics of these attacks deserve a closer look. Backdoor or C2 functionality shows up in 35% of malware-driven breaches, giving attackers a handy way to maintain access, recon the environment and perpetrate more illicit operations long after the initial infection.

The primary vector of infection is via Web applications (Figure 83), which serve as the front door in 71% of cases. To illustrate, we need only look back to the late summer of 2025. A well-known ransomware gang – the same group behind the 2023 MOVEit exploitation – shifted their sights to a zero-day vulnerability in Oracle's E-Business Suite.[107] This campaign resulted in more than 100 organizations being compromised and subjected to extortion, with a heavy concentration of those victims residing right here in the Education sector.

Meanwhile, web application downloads are close behind at 65%, and email attachments are doing their usual damage in 52% of incidents.

Social attacks appear in 22% of breaches and, no surprise here, they're mostly classic phishing attacks (81%), with Email acting as the primary delivery vehicle in 88% of those cases. Meanwhile, Errors are less common in this vertical, showing up in just 17% of breaches. Misdelivery remains the leading error variety, consistent with last year's findings. However, this year, the picture shifts slightly with Loss now accounting for 21% of errors, edging out Misconfiguration, which played a more prominent role in the prior report.

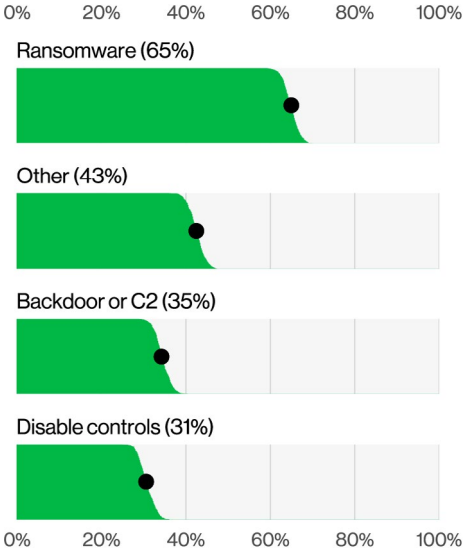

Figure 82. Top Malware action varieties in Educational Services breaches (n=631)

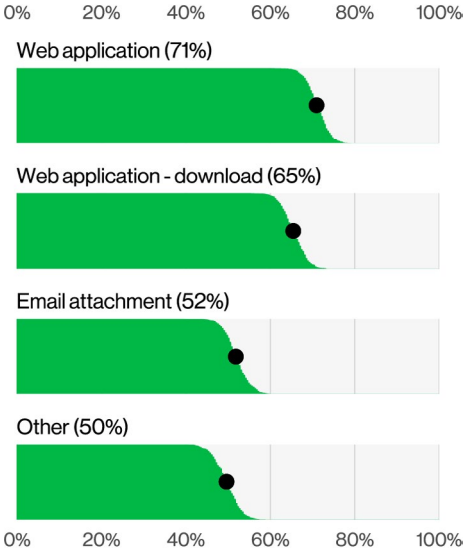

Figure 83. Top Malware vectors in Educational Services breaches (n=272)

106. Ok, already know for sure
107. github.com/vz-risk/VCDB/issues/22574

Financial and Insurance NAICS 52

Summary

This sector continues to be heavily targeted by financially motivated external attackers, with Ransomware-driven System Intrusions, Phishing, Exploit vulnerability and Use of stolen credentials being the primary threats. Human error and third-party exposure remain significant contributing factors.

What is the same?

System Intrusion remains the top pattern since 2022. Attackers remain primarily financially motivated.

Frequency	3,809 incidents, 1,300 with confirmed data disclosure
Top patterns	System Intrusion, Social Engineering and Everything Else represent 81% of breaches
Threat actors	External (88%), Internal (12%) (breaches)
Actor motives	Financial (98%), Espionage (3%) (breaches)
Data compromised	Internal (53%), Personal (43%), Other (28%), Credentials (26%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (22%), Phishing (20%), Credential abuse (15%) (breaches)
Other metrics	Human element (65%), Third-party (34%)

It's a rich man's world.

This sector continues to be a favorite among attackers, which isn't surprising given that its core business is handling money. Incidents increased this year, but the proportion that resulted in successful breaches remains fairly consistent.

As Figure 84 shows, the System Intrusion pattern has been on a strong upward trajectory over the past several years. After briefly competing with Social Engineering in the 2022 and 2023 DBIRs, it has reigned supreme ever since. If you're not familiar with our attack patterns, System Intrusion is the more complex pattern, where the attacks that take a bit more work tend to live. It has been dominated by Ransomware, and we also see the Use of stolen creds and Exploit vuln hacking actions appearing frequently as the primary methods of initial access.

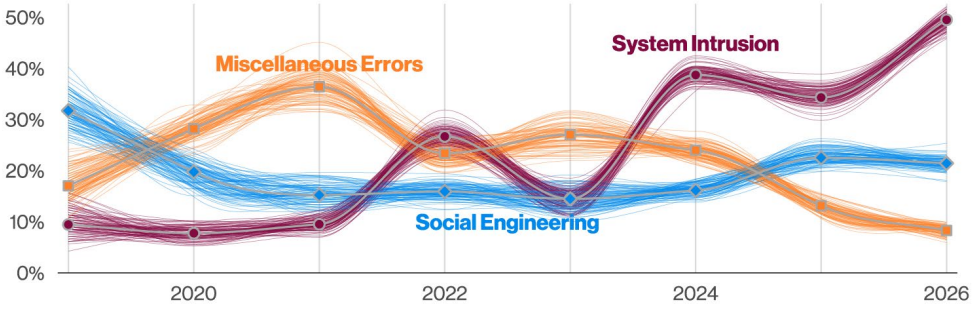

Figure 84. Top patterns in Financial and Insurance breaches over time (n for 2026 dataset=1,300)

The Social Engineering pattern has long played second fiddle to System Intrusion, and this year is no different. The top social actions in this sector were Phishing and Pretexting. Phishing, being the lower-effort attack, is seen more than twice as often as Pretexting. Criminals are nothing if not efficient in their efforts to gain access to your systems, and we saw significant examples of compromises targeting this industry being initiated by social engineering attacks against third parties this past year.

This year, the Everything Else pattern has moved into third place. This isn't a true pattern but a catch-all for cases that lack enough detail to be classified elsewhere. Its rise to third place for this sector is notable, however, since Basic Web Application Attacks held that spot in the 2025 report. This may turn out to be an interesting change, but it may also be due to the normal year-to-year data variation from our partners.

Attackers in this industry are still predominantly External and financially motivated. Internal actors decreased substantially since last year, from 22% to just 12%. However, it is important to remember that Internal actor breaches are mainly accidental and not carried out with malicious intent.

Healthcare NAICS 62

Summary

Healthcare organizations face a mix of Ransomware-driven System Intrusions and persistent human errors, with financially motivated external attackers exploiting vulnerabilities, Phishing, and using stolen credentials. Staff mistakes and Misconfigurations remain a chronic source of breaches.

What is the same?

Miscellaneous Errors has been in the top patterns for this industry since we started keeping track. System Intrusion is, once again, in the number one spot for the second year in a row.

Frequency	1,492 incidents, 1,438 with confirmed data disclosure
Top patterns	System Intrusion, Miscellaneous Errors and Social Engineering represent 81% of breaches
Threat actors	External (81%), Internal (19%) (breaches)
Actor motives	Financial (99%), Espionage (2%) (breaches)
Data compromised	Internal (65%), Personal (37%), Credentials (25%), Other (19%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (20%), Phishing (14%), Credential abuse (11%) (breaches)
Other metrics	Human element (54%), Third-party (32%) (breaches)

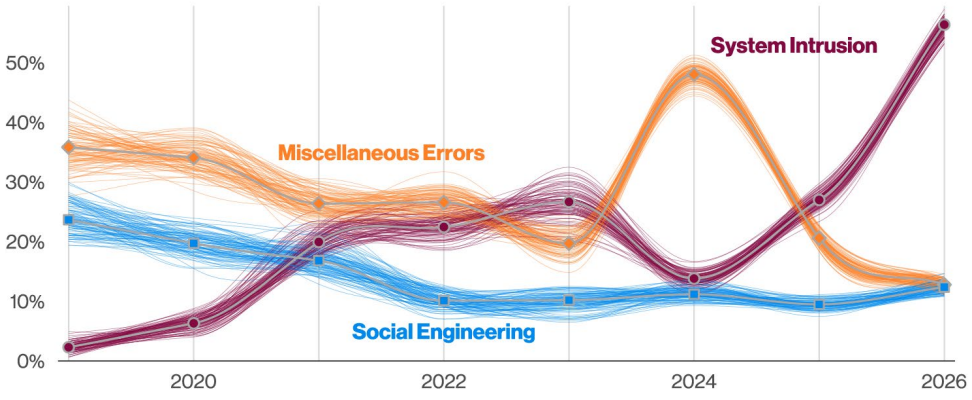

Figure 85. Top patterns in Healthcare breaches over time (n for 2026 dataset=1,438)

Snake oil remedies come and go on social media, often promising "this one weird trick" will melt away belly fat with no effort on your part. These claims can be ignored, but basic hygiene, whether personal or cyber, cannot. The fundamental principles must be addressed for an organization to be able to weather cyber incidents and breaches.

Breaches take many forms, but in the Healthcare sector, one pattern stands out: Miscellaneous Errors. DBIRs from 2014 through 2026 have shown that Healthcare has been among the most affected by staff mistakes. Miscellaneous Errors has been among the top three patterns each year. The ranking may vary from year to year, but it remains a chronic problem that needs a cure.

This year's top errors in Healthcare were Misdelivery (data is delivered to the wrong recipient, in any format) Loss (often involving unencrypted user devices and portable media) and Misconfiguration (such as exposing a data store to the internet without appropriate controls). These Misconfigurations are frequently discovered by security researchers who typically make an effort to notify the victim organizations rather than simply take the data for their own use.

Figure 86 illustrates that despite repeated recommendations to implement controls to prevent or limit the impact of such mistakes, these Errors have remained persistent over the years. Controls to combat these kinds of mistakes, which appear in our dataset again and again, would be part of those security fundamentals we mentioned.

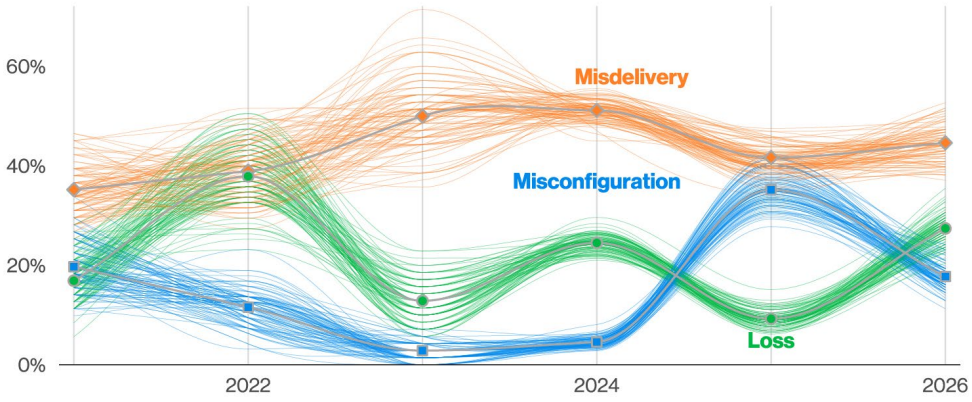

Figure 86. Top Error varieties in Healthcare breaches over time (n for 2026 dataset=186)

Hopping off our soapbox

System Intrusion remains the top pattern for the Healthcare industry and is largely driven by Ransomware. Threat actors commonly gain access via the Use of stolen credentials or by the Exploitation of vulnerabilities, then deploy Ransomware and frequently follow up with the exfiltration of data for future leverage. In other cases, Actors simply rely on extortion alone and do not trigger the encryption of data.

This year, many industries were affected by the Oracle E-Business Suite vulnerability,[108] and Healthcare was not spared. A number of organizations reported having been affected by this attack, largely attributed to the Cl0p criminal group. In part, this contributed to the third-party figure of 32% of breaches in this sector and illustrates how the complex web of third-party relationships can impact any organization.

The security fundamentals do not only apply to your own organization – they must also be baked into the contracts made with business associates and suppliers, as well.

The Social Engineering pattern re-entered the top three, replacing last year's Everything Else. For social actions, Phishing was most common, followed by Pretexting. This year, we saw a small number of Baiting cases in which attackers infected the websites frequently visited by their targets.

108. z2data.com/insights/everything-you-need-to-know-about-the-oracle-data-breach

Manufacturing NAICS 31–33

Summary

The number of breaches in this industry continues to grow, with the uptick in numbers largely due to Ransomware attacks.

What is the same?

The top patterns for Manufacturing are the same as last year, with the vast majority of actors being financially motivated.

Frequency	3,627 incidents, 2,713 with confirmed data disclosure
Top patterns	System Intrusion, Social Engineering and Basic Web Application Attacks represent 91% of breaches
Threat actors	External (95%), Internal (5%) (breaches)
Actor motives	Financial (87%), Espionage (15%) (breaches)
Data compromised	Internal (81%), Credentials (26%), Other (22%), Personal (17%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (38%), Phishing (13%), Credential abuse (11%) (breaches)
Other metrics	Third-party (61%), Human element (56%) (breaches)

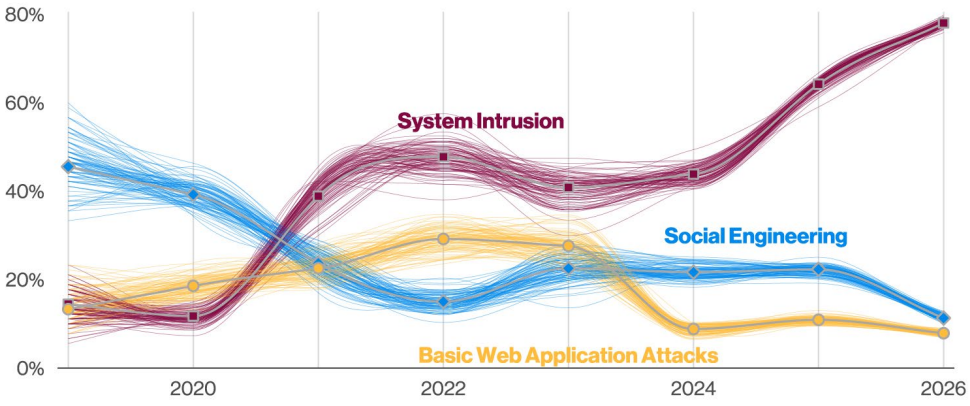

Figure 87. Top patterns in Manufacturing breaches over time (n for 2026 dataset=2,713)

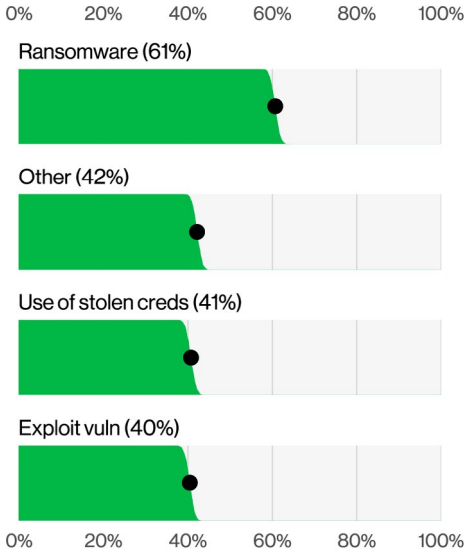

Figure 88. Top Action varieties in Manufacturing breaches (n=2,417)

Ransomware is still, in large part, the driving force behind both the growth in breaches and the prominence of System Intrusion incidents. Malware was involved in 75% of the breaches in this vertical, with Ransomware accounting for 61% (Figure 88). A prominent example within this sector is the late 2025 Ransomware attack on the Japanese company Asahi Group Holdings.[109] The incident forced a shutdown of their domestic manufacturing facilities and resulted in a suspension of shipments, while also potentially compromising corporate data. This event illustrates that the financial impact of a breach can often extend far beyond the immediate ransom or extortion demands, as the operational downtime and downstream supply chain disruption can be considerable.

109. github.com/vz-risk/VCDB/issues/22512

Hacking actions were involved in 71% of Manufacturing breaches. The main tactics haven't changed much since last year's report: Use of stolen credentials and Exploit vulnerability each contributed to 41% of Manufacturing breaches.

Social Engineering may be the second most common pattern in this vertical, but its actions still lag well behind the leaders Malware and Hacking, as shown in Figure 89. Social actions appeared in only 16% of breaches, and most of those were of the Phishing variety (77%). The more elaborate Pretexting schemes barely register by comparison, accounting for only 18% of social attacks in Manufacturing breaches.

Internal data in this sector apparently didn't get the memo about staying put. It's involved in 80% of breaches (see Figure 90), making emails, plans and reports among the favorite items on the criminal's takeout menu, or perhaps simply the easiest to obtain. Credentials show up next in 26% of incidents, and Personal (personally identifiable information, or PII) data makes an appearance in 17% of Manufacturing breaches – less common but still more than enough to ruin someone's day.

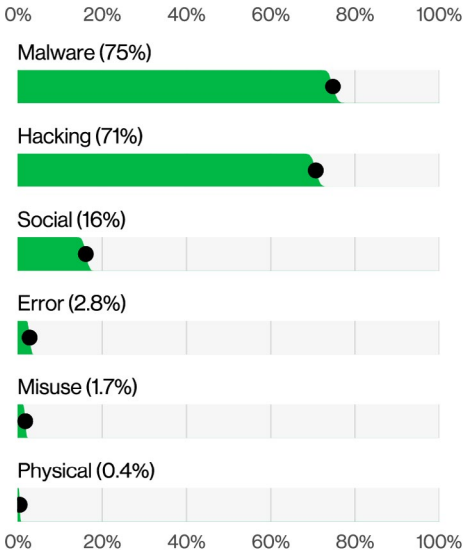

Figure 89. Action types in Manufacturing breaches (n=2,713)

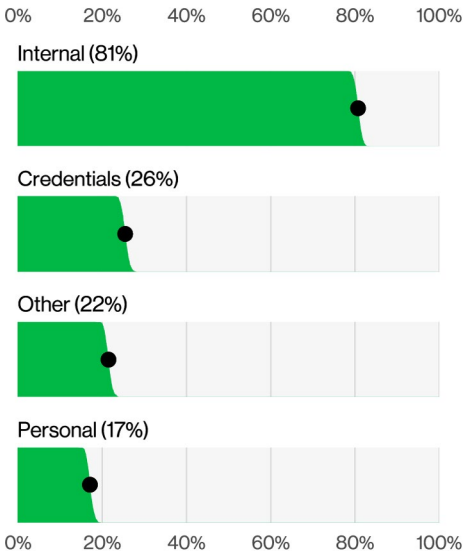

Figure 90. Top Data varieties in Manufacturing breaches (n=2,452)

Public Administration NAICS 92

Summary

Public Administration is primarily targeted by a combination of financially motivated criminals and State-affiliated actors, leading to a high frequency of System Intrusion via vulnerability exploitation and Ransomware. Additionally, the sector faces an unusually high rate of internal incidents driven by Miscellaneous Errors – specifically Misdelivery due to the sheer volume of correspondence – as well as intentional data mishandling.

What is the same?

The first two attack patterns in this sector remain the same as last year, although Basic Web Application Attacks gave way to Privilege Misuse this year. The prevalence of External actors targeting this industry remains consistent YoY.

Frequency	3,634 incidents, 2,410 with confirmed data disclosure
Top patterns	System Intrusion, Miscellaneous Errors and Privilege Misuse represent 80% of breaches
Threat actors	External (56%), Internal (44%) (breaches)
Actor motives	Financial (69%), Espionage (33%), Ideology (2%) (breaches)
Data compromised	Personal (50%), Internal (39%), Other (37%), Secrets (30%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (40%), Phishing (20%), Credential abuse (8%) (breaches)
Other metrics	Human element (69%), Third-party (36%)

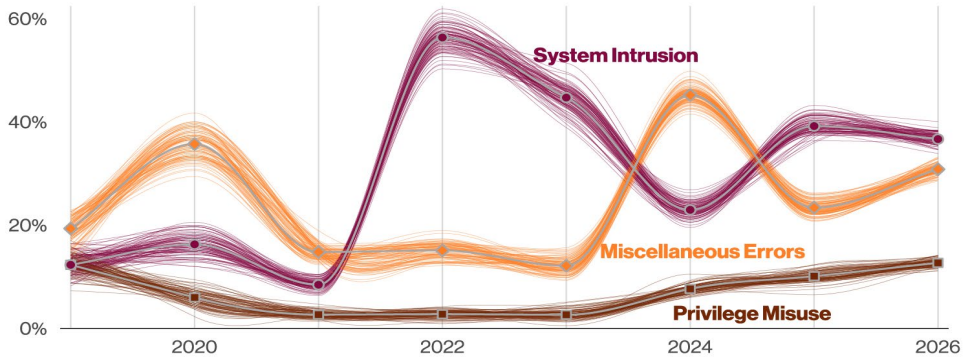

Figure 91. Top patterns in Public Administration breaches over time (n for 2026 dataset=2,410)

This year, the top three incident patterns in Public Administration are System Intrusion, Miscellaneous Errors and Privilege Misuse (Figure 91). It's worth noting, however, that Privilege Misuse is only 0.01% more common than Social Engineering, so those two are still essentially battling for supremacy.

At first glance, the prominence of Error and Misuse might make you think, "Is the government really this prone to mistakes and bad behavior?" Possibly – but there are some important context points to keep in mind when looking at NAICS 92:

- We receive incident data about government entities from a limited number of contributors, and that number is much smaller than the contributors reporting on private sector incidents. Because the Public Administration dataset is smaller, the results are more susceptible to bias.
- On the positive side, the government-related data we do receive tends to be detailed and high quality. This allows for a more granular, in-depth view of what's happening, even though it reflects fewer organizations overall.
- An even more critical factor is that government entities often operate under stricter regulatory and reporting requirements than most private sector organizations. They are often required to report more types of incidents and at a higher level of detail.

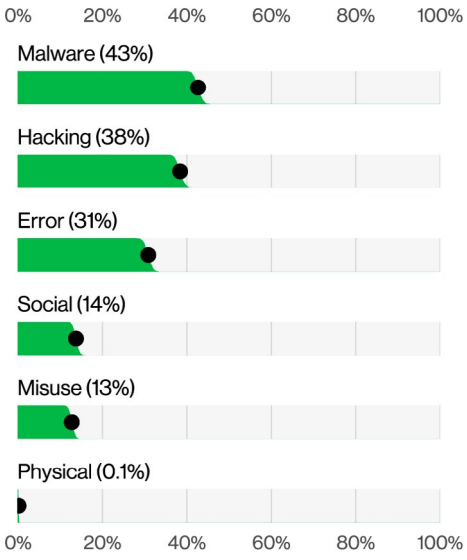

Figure 92. Top Actions in Public Administration breaches (n=2,410)

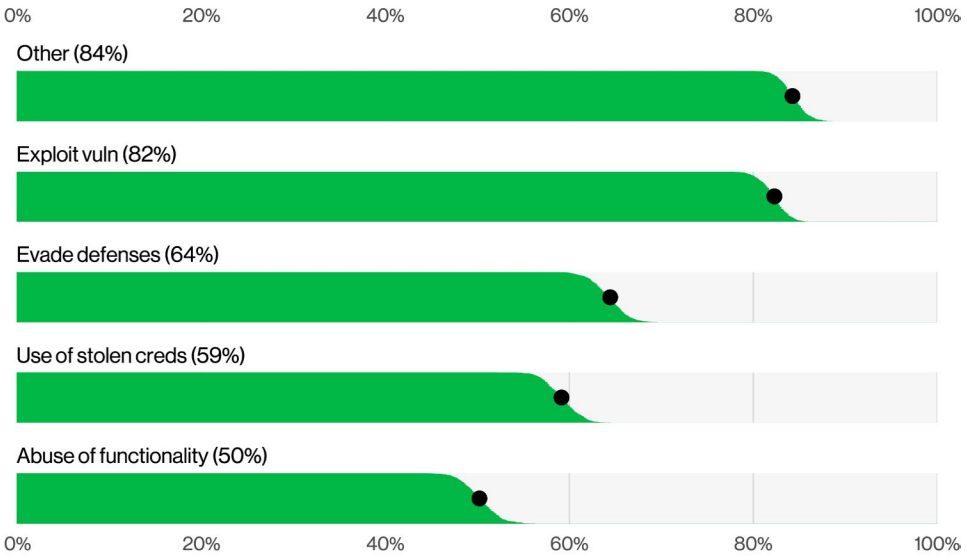

Figure 93. Top Hacking varieties in Public Administration breaches (n = 661)

With those caveats in mind, we can now look more closely at what drives the patterns that appear at the top of the list.

In Public Administration, Hacking (38%) and Malware (43%) actions show up in roughly equal proportions (Figure 92), which tracks with what we see across most industries this year. Given that more complex attacks tend to blend multiple techniques (and with System Intrusion sitting in the top spot), that balance is not exactly shocking. What is notable, however, is that Hacking and Malware appear at lower rates in breaches here than in many other verticals.

When we examine hacking in greater detail (Figure 93), the Exploitation of vulnerabilities is predominant, accounting for 82% of hacking-related breaches in government. Evade defenses also appears at an unusually high rate, present in 64% of breaches involving hacking. Finally, the Use of stolen credentials is a common tactic for threat actors targeting government entities, occurring in 59% of hacking-related breaches.

Malware actions such as Ransomware, Backdoors, C2, disabling controls and evading defenses are in frequent use. If we require any further clarity on why the patterns fall where they do, a look at these Malware actions provides it. This blend of activity reflects how Ransomware and other complex attack scenarios typically unfold.

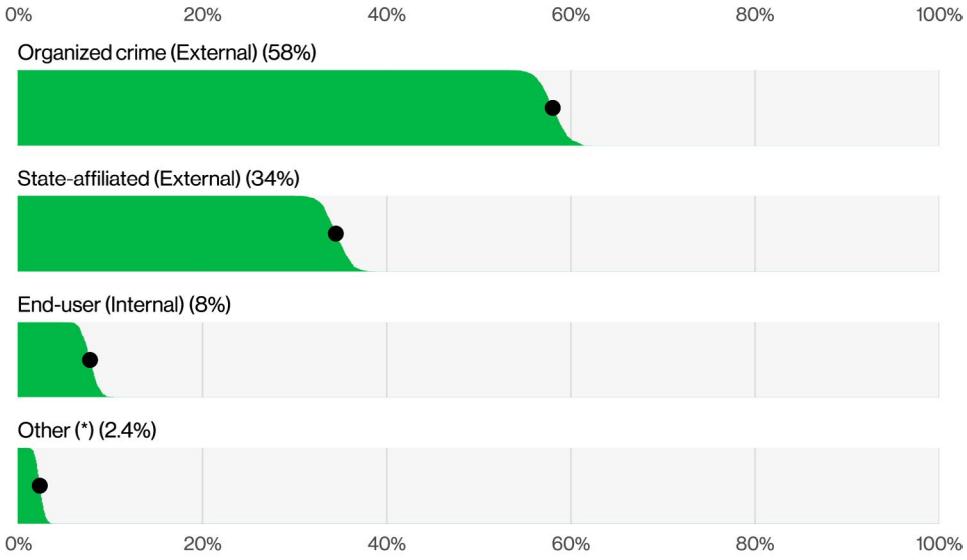

Figure 94. Actor varieties in Public Administration breaches (n=1,146)

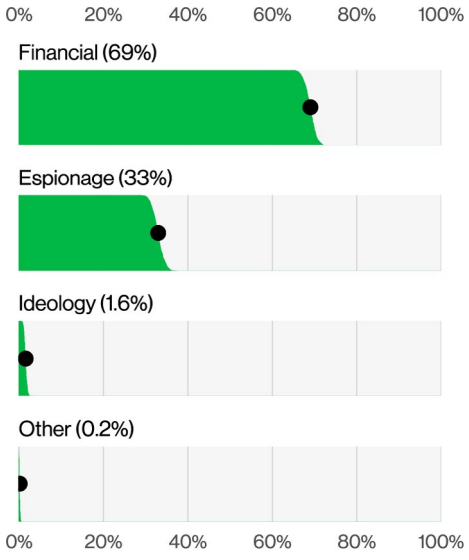

Figure 95. Actor motives in Public Administration breaches (n=1,201)

In this segment, a quick look at the main threat actors is quite illuminating. Figure 94 does a nice job of summarizing who is driving the action in Public Administration. External actors make up a slim majority of breaches at 56%, but Internal actors still accounted for a substantial portion of incidents in NAICS 92.

On the External side, organized criminal groups are the primary offenders and are, unsurprisingly, mostly in it for the money (69% financially motivated).

State-affiliated actors also feature heavily, appearing in just over one-third of breaches (35%) and frequently acting with an Espionage motive (33%) (Figure 95). One example of this kind of activity was the breach of the U.S. Department of the Treasury by the Chinese nation state hacking group Silk Typhoon, in which a vulnerability present in the software of a third party's cloud-based support services was exploited.[110]

This type of action is part of a broader pattern of Chinese cyberespionage aimed at U.S. critical infrastructure and national security frameworks primarily due to rising global political tensions.

In other words, this sector is contending with a mix of financially motivated criminal groups and state-aligned actors focused on intelligence gathering. This is what ultimately puts System Intrusion in the number one spot among the patterns.

110. govinfosecurity.com/report-chinese-hackers-breached-cfius-a-27274?rf=2025-01-13_ENEWS_SUB_GIS__Slot1_ART27274

We all make (the same) mistakes.

The Miscellaneous Errors pattern makes a strong showing in this sector: Instances of this pattern (31%) are considerably higher here than in most other verticals. Misdelivery, once again, rises to the top as the dominant error type, keeping pace with what we see in nearly every other sector. Misdelivery is, at its core, what it sounds like: Data intended for one recipient is sent to someone else. Of all the industries where this can happen, Public Administration is perhaps the easiest to understand. There is a large chain of individuals to notify – often by letter or email – that they did something wrong or that they owe a few dollars more to the powers that be, and that sheer volume creates room for mistakes.

Misdelivery can also involve printed correspondence or electronic communications. Whether it is a physical letter, an email or a form that must be filled out in triplicate and returned, it is surprisingly easy to make an "off by one" error and send it to the wrong person. Misdelivery accounts for 88% of all errors in the Public Sector. For comparison, the second most common error type – Classification error – comes in at just 4%. If your organization sends out a large volume of correspondence, it would be wise to implement safeguards to reduce this all-too-common problem (and the government is far from alone here; large healthcare and insurance firms with their constant mailings face similar exposure).

For anyone hoping for a silver lining,[111] it's worth noting that most government errors are not the result of inadequate processes (1%) or poor technology (9%) but of plain old Carelessness (91%) – which does not exactly inspire confidence.

You did that on purpose!

Shifting from honest mistakes to intentional behavior, misuse-related breaches in this sector are overwhelmingly about Data mishandling (82%). At its core, Data mishandling is the inappropriate use of data, and it wears many hats: Sending information through unauthorized channels, convenient workarounds that put data at risk, or storing data in ways that don't meet policy or regulatory requirements all qualify.

The second most common flavor of misuse is Privilege abuse, appearing in 18% of misuse-related breaches. In these cases, the actors are not breaking in – they are simply logging on. They deliberately use the legitimate access they already have to systems or data for an improper purpose, typically to gain some real (or at least perceived) personal benefit.

Misuse can be insidious and difficult to detect if all of your organization's controls are outward facing, looking only for the attacker trying to get in from the outside. Part of those security fundamentals is to make sure you can detect malicious activity, even when it is "coming from inside the house."

111. We will be sure to let you know as soon as we find one.

Retail NAICS 44–45

Summary

Retail organizations face persistent threats from external attackers exploiting vulnerabilities, stealing credentials and Phishing. These activities often lead to ransomware attacks and data theft, with third-party systems and internal corporate data becoming increasingly valuable targets.

What is the same?

The top three patterns remained the same, but their order of supremacy shifted a bit. The patterns have been the same consistently for many years now, but which is more prevalent in a given year changes.

Frequency	997 incidents, 806 with confirmed data disclosure
Top patterns	System Intrusion, Basic Web Application Attacks and Social Engineering represent 95% of breaches
Threat actors	External (99%), Internal (1%) (breaches)
Actor motives	Financial (85%), Espionage (19%) (breaches)
Data compromised	Internal (84%), Credentials (26%), Secrets (20%), Other (14%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (42%), Credential abuse (14%), Phishing (9%) (breaches)
Other metrics	Third-party (68%), Human element (58%) (breaches)

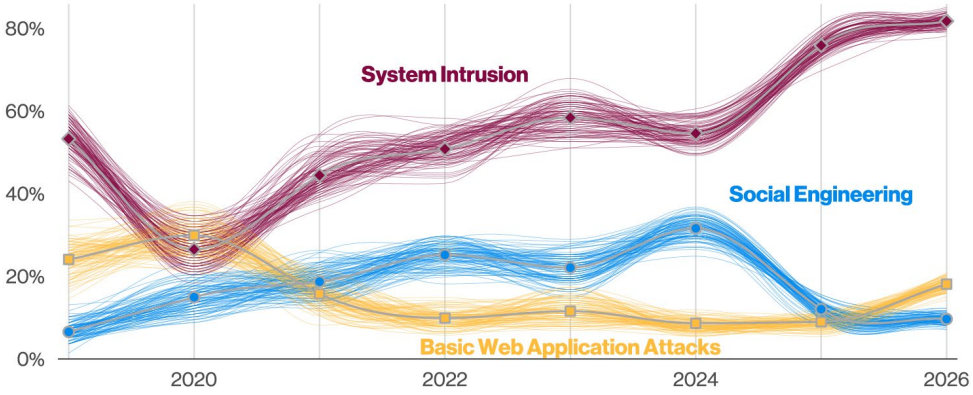

Figure 96. Top patterns in Retail breaches over time (n for 2026 dataset=806)

I have a coupon code.

Few shoppers can resist a good deal. This is also true of attackers, and they have been compromising systems like they were on final clearance.

This year, the number of incidents rose slightly, but the number of breaches nearly doubled. Despite that, the top three patterns essentially remain the same as last year, just in a different order. System Intrusion still leads, while Basic Web Application Attacks and Social Engineering continue to dance around each other.

As Figure 96 illustrates, it is not until you look all the way back to the 2020 DBIR that you see a change in the membership of the top three patterns. While they have shuffled about a bit on the stage, they have been close compatriots for several years. This consistency tells us that the same kinds of attacks are often being leveraged against this industry's infrastructure year after year – with a certain level of success. Case in point, this past year saw clothing retailer Hot Topic experiencing a breach affecting 57 million customers.[112] Clearly there remains significant incentive for attackers to target this sector, given the sheer volume of data up for grabs.

The unholy trio of Ransomware, Exploit vuln and Use of stolen creds figured prominently in the actions taken in this sector (Figure 97), and when combined with social attacks, you have accounted for the most prevalent ways in. Ransomware remains an ongoing problem across this industry and was the top malware action in breaches in this sector.

112. github.com/vz-risk/VCDB/issues/21205

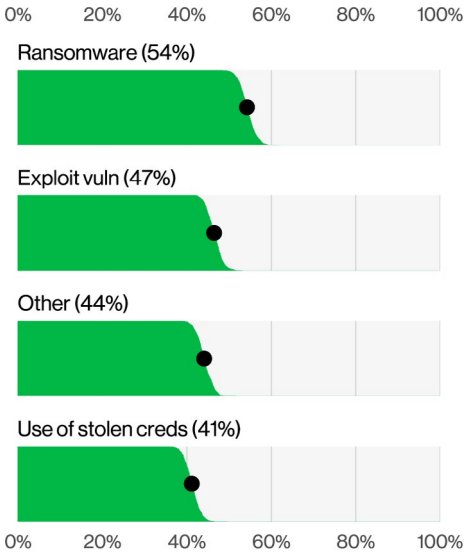

Figure 97. Top Action varieties in Retail breaches (n=719)

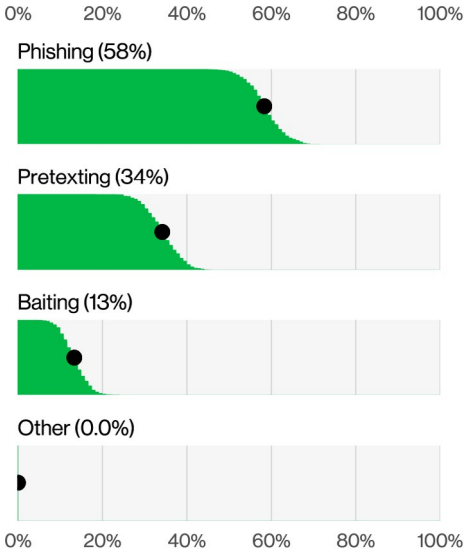

Figure 98. Top Social actions in Retail breaches (n=120)

The social varieties most commonly seen in these breaches have been Phishing and Pretexting, with the former almost twice as common as the latter, as seen in Figure 98. That makes some sense, since Phishing tends to be the lower-effort attack, where Pretexting tends to take a bit more time and skill to achieve a successful result (from the attacker's perspective). However, since both tactics have been quite successful for the attackers, it behooves organizations to have easy methods for people to report when they have become victims of these kinds of attacks.

Espionage-motivated actors increased again this year, rising from 9% to 19% of breaches where the motive was known. This suggests more sophisticated actors have taken notice of this sector and are turning their attention to what kinds of useful data their victims may have.

While this sector once saw primarily Payment card data compromised, threat actors have evolved and now target any data they can monetize, leading to a more diverse mix of data types being affected. Internal data, which can include plans, strategies and other information of value to Espionage-motivated attackers and ransomware actors looking for leverage rose from 65% last year to 84%. Figure 99 has the details.

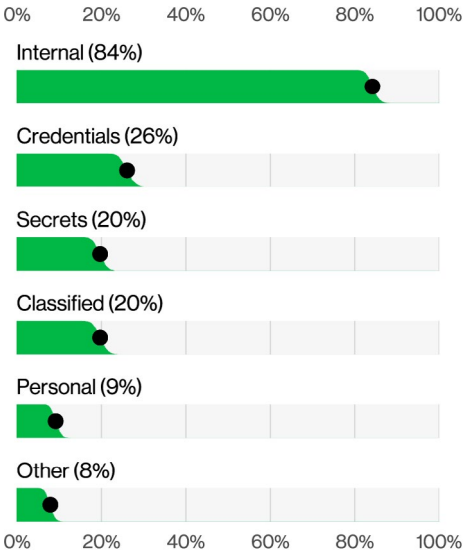

Figure 99. Top Data varieties compromised in Retail breaches (n=721)

Focused analysis

/06

Small- and medium-sized businesses

Summary

Small organizations are disproportionally impacted by Ransomware and face many of the same threats as other industries and organizations but often with less resources available.

What is the same?

System Intrusion, Basic Web Application Attacks and Social Engineering continue to be the main drivers of breaches in SMBs.

Frequency	7,256 incidents, 7,152 with confirmed data disclosure
Top patterns	System Intrusion, Basic Web Application Attacks and Social Engineering represent 100% of breaches
Threat actors	External (100%) (breaches)
Actor motives	Financial (100%) (breaches)
Data compromised	Internal (97%), Credentials (31%), System (1%), Other (1%) (breaches)
Initial access vector breakdown	Exploitation of vulnerabilities (26%), Credential abuse (13%), Phishing (9%) (breaches)
Other metrics	Third-party (55%), Human element (45%) (breaches)

Here, we view our data through the lens of organizational size, to give a perspective for the smaller organizations that can get lost in the overall dataset. For our purposes, organizations with fewer than 1,000 employees are considered small businesses. Hiring that 1,001st employee moves you to the enterprise category. Welcome to the big leagues!

Being a small organization, you may mistakenly think that your threat profile and who would be interested in compromising you are significantly different from everyone else. That is, until you find yourself on the wrong side of a ransomware attack. Overall, SMBs face similar types of threats as everyone else, including the same breach patterns that show across many different industries, and this has been the case for many years now. While the top three patterns changed slightly (Figure 100), System Intrusion remains the top pattern in small organization breaches, while Social Engineering and Basic Web Application Attacks switched positions. Unsurprisingly, financially motivated External actors are perpetrating the majority of the breaches here.

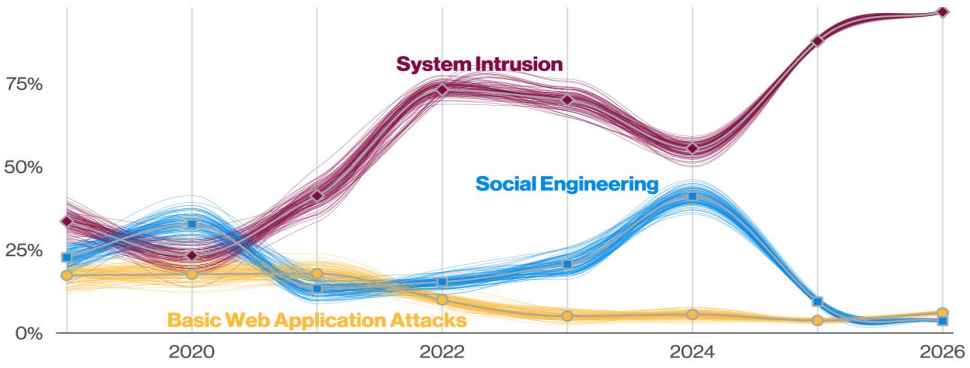

Figure 100. Top patterns in SMB breaches over time (n for 2026 dataset=7,152)

When we examine the specific actions impacting SMBs, we're struck with a heavy dose of déjà vu, as we see Ransomware, Use of stolen credentials and Exploit vulnerabilities show up in the top (Figure 101). As we've discussed in our section on Ransomware (found in the "System Intrusion" pattern section), these actors are often casting out wide nets to compromise as many organizations as possible in the hope that at least some of the victims will pay. For many of these attacks, they're opportunistic in nature, and it's not so much about industry or the revenue of the victims but the fact that the victims had credentials that were compromised (38%) or unpatched vulnerabilities in edge devices (29%) that resulted in them being victimized. Of the Ransomware cases where we have information on the organization size, we found that about 96% of Ransomware victims were SMBs. While SMB Ransomware cases may rarely make the news, they certainly make it into our dataset.

In terms of the data stolen from SMBs, Internal and Credentials remain the primary data types taken, while Personal dropped off the list entirely. Internal data is one of the common enumerations that's selected when it comes down to these ransomware cases, as it's not always clear as to what was taken, but odds are it was likely some type of non-public data that the actors are then trying to extort.

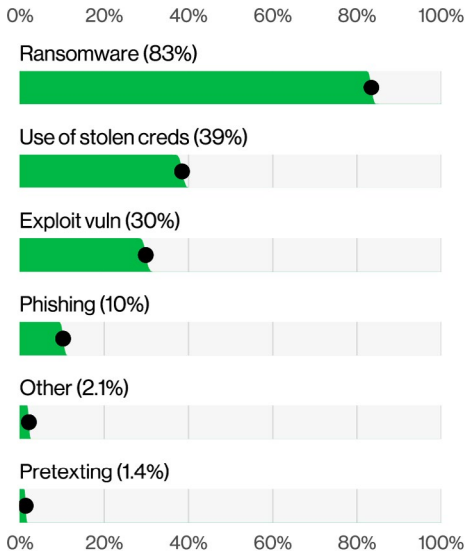

Figure 101. Top Action varieties in SMB breaches (n=6,182)

Regions

/07

Regional analysis

In this section, we examine breaches from a macro-regional perspective to highlight how trends differ or remain consistent across geographical regions.

Our visibility into any given area is determined by several variables, including regional disclosure laws, our specific datasets and the locations where our contributors conduct business. If you feel your region is not adequately represented in the following pages, please contact us about becoming a data contributor and encourage other organizations in your area to do the same.

We define world regions in accordance with the United Nations M49[113] standards, which combine super-regions and sub-regions. Based on this framework, we will examine the following regions:

APAC: Asia and the Pacific, including Southern Asia (034), South-eastern Asia (035), Central Asia (143), Eastern Asia (030) and Oceania (009)

EMEA: Europe, Middle East and Africa, including North Africa (015), Europe (150) and Eastern Europe (151) and Western Asia (145)

LAC: Latin America and Caribbean, which consists of breaches in South America (005), Central America (013) and Caribbean (029)

NA: Northern America (021), which primarily consists of breaches in the United States and Canada

Longtime readers will recognize the at-a-glance tables located at the top of each major section. We have combined these to provide a quick comparison across all regions regarding incident frequency, top patterns and other key metrics.

Region	Frequency	Top patterns	Threat actors	Actor motives	Data compromised	Initial access vectors	Misc
APAC	5,229 incidents, 2,855 with confirmed data disclosure	System Intrusion, Basic Web Application Attacks and Social Engineering represent 97% of breaches	External (99%), Internal (1%) (breaches)	Financial (70%), Espionage (36%) (breaches)	Internal (70%), Credentials (36%), Other (35%), Secrets (30%) (breaches)	Exploitation of vulnerabilities (42%), Credential abuse (25%), Phishing (15%) (breaches)	Third-party (69%), Human element (71%) (breaches)
EMEA	8,245 incidents, 6,060 with confirmed data disclosure	System Intrusion, Social Engineering and Miscellaneous Errors represent 92% of breaches	External (80%), Internal (20%) (breaches)	Financial (76%), Espionage (27%) (breaches)	Internal (73%), Other (49%), Personal (34%), Secrets (24%) (breaches)	Exploitation of vulnerabilities (47%), Phishing (28%), Credential abuse (6%) (breaches)	Third-party (54%), Human element (70%) (breaches)
LAC	813 incidents, 718 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 98% of breaches	External (99%), Internal (1%) (breaches)	Financial (90%), Espionage (11%) (breaches)	Internal (93%), Credentials (23%), Secrets (24%), Other (3%) (breaches)	Exploitation of vulnerabilities (44%), Phishing (20%), Credential abuse (5%) (breaches)	Third-party (74%), Human element (57%) (breaches)
NA	12,371 incidents, 8,426 with confirmed data disclosure	System Intrusion, Social Engineering and Basic Web Application Attacks represent 87% of breaches	External (88%), Internal (12%) (breaches)	Financial (98%), Espionage (3%) (breaches)	Internal (77%), Credentials (36%), Personal (9%), Other (8%) (breaches)	Exploitation of vulnerabilities (30%), Credential abuse (20%), Phishing (12%) (breaches)	Third-party (43%), Human element (59%) (breaches)

Table 5. At-a-glance table by region

113. unstats.un.org/unsd/methodology/m49

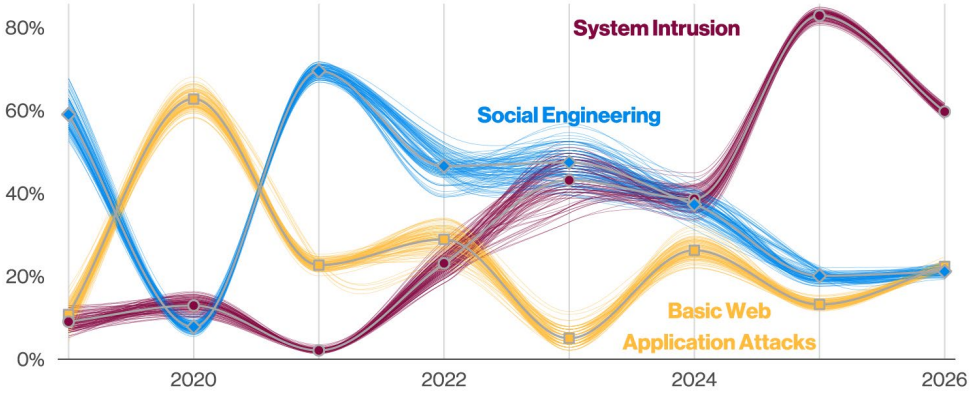

Figure 102. Top patterns in APAC breaches over time (n for 2026 dataset=2,855)

The APAC region

APAC is still contending with the same familiar patterns that have topped the charts there for the last few years. System Intrusion continues to lead by a wide margin and is responsible for 60% of breaches – still roughly three times the share of the next closest patterns, even though it experienced a sharp drop from 89% in the 2025 report. Social Engineering remains fairly steady with last year's level at 21%. The standout change comes from Basic Web Application Attacks, which have doubled since last year and now account for 22% of breaches. That rise likely correlates with the prominence of the Use of stolen credentials as the leading Hacking action, since those are often the fuel that powers the shorter, less-sophisticated Basic Web Application Attacks-style attacks.

APAC shares the same leading patterns – and by extension, the same leading Action types – as the overall dataset, but some of them show up more frequently here than elsewhere.

Hacking is involved in 83% of breaches in APAC and Malware in 71%, compared to 64% and 63%, respectively, in the overall data, which is quite a noticeable jump (Figure 103). Some of that difference can be chalked up to contributor bias in the region, based on who sends us data and what kinds of cases they see. Nevertheless, that combination of Hacking and Malware is a hallmark of multistep, more-complex breaches that require both lateral movement and persistence. Therefore, a significant portion of this uplift likely reflects the higher volume of Espionage-related attacks in the APAC region.

The topic of espionage provides us with an excellent entry point for discussing who's doing the attacking in APAC and what's driving them. Threat actors here are almost entirely External, accounting for 99% of breaches. State-affiliated actors are responsible for a striking 36% of breaches – more than in any other region. Again, a certain amount of that is a reflection of our contributor base in APAC, and some of it also reflects regional geopolitics and where these targets sit on the map.

Nevertheless, financially motivated organized crime groups still account for the bulk of breaches in APAC and were involved in 67%. The influence of organized crime groups in this region can clearly be seen in the July 2025 breach affecting the airline Qantas.[114] The personal data of more than five million customers was stolen by a group known as Scattered Lapsus$ Hunters via a third-party platform. The criminal group then placed extortion demands on the victim and released the data when the ransom was not paid. This incident was one of Australia's largest since 2022 and underscores the continued risk of third-party data stores and the downstream impact they can create.

With regard to the types of data stolen, in case our readers have any lingering doubts regarding how strong the Espionage signal is in APAC, the data varieties should put them to rest. Secrets appear in 28% of APAC breaches. In the overall dataset, those numbers drop to 13%. That's a clear indication that APAC is experiencing more breaches, in proportion where highly sensitive information is the main prize.

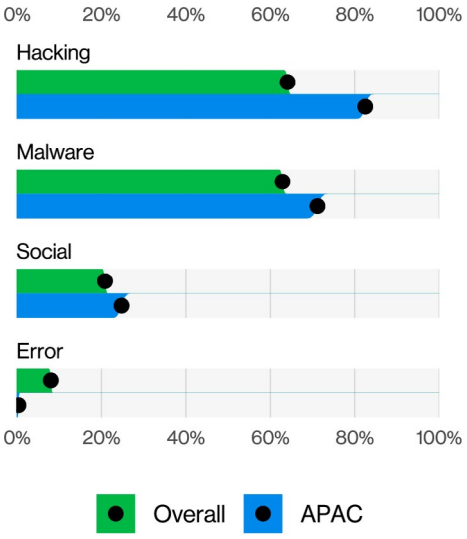

Figure 103. Top Action varieties in APAC breaches (n for APAC=2,855)

114. reuters.com/sustainability/boards-policy-regulation/qantas-says-customer-data-released-by-cyber-criminals-months-after-cyber-breach-2025-10-12

Collective Cyber Resilience in Action: Operational Lessons from Responding to UNC3886

David Koh

Commissioner of Cybersecurity and Chief Executive of the Cyber Security Agency (CSA) of Singapore

Cyberspace is a shared domain, and securing it is a team sport. Sustaining digital trust therefore relies on collective resilience, rather than individual measures alone. The campaign by Advanced Persistent Threat (APT) actor UNC3886 against Singapore's telecommunications infrastructure was a timely reminder of this reality. The sophistication of the campaign made detection challenging, and necessitated a coordinated national response across government and industry. From this experience, we share several operational reflections that may be useful to fellow cyber defenders.

First, early detection and trusted reporting channels are crucial. The threat actor's activities were first detected by our telecommunications operators (telcos), who proactively notified the Singapore authorities. The prompt notification enabled a swift whole-of-Government response, mounted in close partnership with the telcos to contain and remediate the breach.

Second, effective inter-agency coordination was central to the response. Codenamed 'Operation Cyber Guardian', the effort became Singapore's largest coordinated cyber incident response to date. More than 100 cyber defenders were mobilised across agencies such as CSA, the Infocomm Media Development Authority (IMDA), the Centre for Strategic Infocomm Technologies (CSIT), the Singapore Armed Forces' Digital and Intelligence Service (DIS), the Government Technology Agency of Singapore (GovTech), and the Internal Security Department (ISD). The breadth and scale of participation reflected the sophistication of the threat, and affirmed coordinated response as a key pillar of Singapore's cyber defence.

Third, public-private partnership is imperative in cybersecurity, which is ultimately a shared responsibility. The telcos proactively cooperated and worked alongside government agencies throughout the operation, supporting investigations and implementing the necessary detection and remediation measures. The telcos have also strengthened their defences through interventions such as joint threat hunting, penetration testing, and uplifting their capabilities. This close partnership enabled effective containment, limited the threat actor's activities, and safeguarded our essential services.

Singapore has taken a transparent and measured approach in communicating the cyber threats that we face. Beyond raising public awareness, transparency also signals resolve and reinforces the collective commitment to safeguard our networks and systems against sophisticated threats. We hope that sharing these lessons will reinforce and demonstrate collective cyber resilience in action, even as the threats grow more complex and frequent.

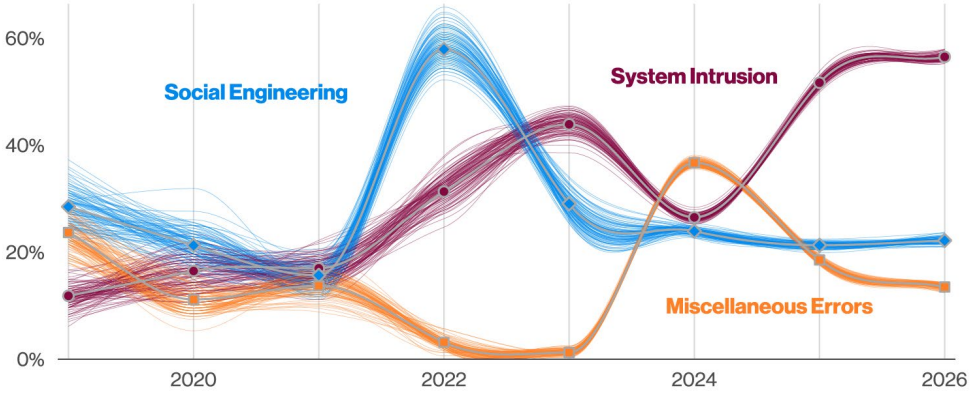

Figure 104. Top patterns in EMEA breaches over time (n for 2026 dataset=6,060)

The EMEA region

System Intrusion, Social Engineering and Miscellaneous Errors are once again headlining the EMEA story, closely mirroring last year. System Intrusion accounts for 57% of breaches in the region, up slightly from 53% last year. Miscellaneous Errors moves in the opposite direction, dropping (though not significantly) from 19% to 14%, while Social Engineering holds steady at 22% of breaches, keeping its usual seat at the table.

In previous reports, we've been up front about a built-in tilt toward North America in our dataset, which is driven largely by where many of our contributors are located. Over the last two to three years, we've pushed to widen that view and bolster coverage across other regions as much as possible. The results suggest that these efforts have been successful to some degree. As the bar chart in Figure 105 shows, EMEA is dealing with many of the same issues we see elsewhere.

For the overall dataset, 63% of breaches involve Malware, while EMEA edges that out ever so slightly with the highest regional share at 66%, up notably from 54% last year. Hacking-related breaches in EMEA are at 59%, a bit lower than the 64% we see across the full dataset. Neither difference is statistically earth-shattering, but together they do suggest that EMEA is now tracking a bit closer to the global picture than it has in previous years. As one might imagine, the high levels of malware and hacking in EMEA are representative of the continued pressure that financially motivated organized crime actors continue to exert by launching frequent Ransomware attacks against targets in this region.

Meanwhile, Social (27%), Error (14%) and Misuse (6%) actions all come in slightly higher here than in the overall dataset.

The most substantive of the differences with regard to EMEA has to do with the Social actions category. Phishing in EMEA shows up in 84% of social-related breaches, which is 15% higher than the overall dataset (69%). This lines up neatly with another EMEA hallmark: a higher share of State-affiliated actors. Across the full dataset, State-affiliated actors are involved in 14% of breaches; in EMEA, that figure jumps to 23%. Not surprisingly, that also tracks with a higher rate of Espionage-motivated breaches in EMEA (27%) compared to the overall data (13%). This is certainly related to the higher percentage of State-affiliated actors that we see here. Considering the complex current political landscape in the region, it doesn't come as a surprise to see the threat of Espionage continue to persist.

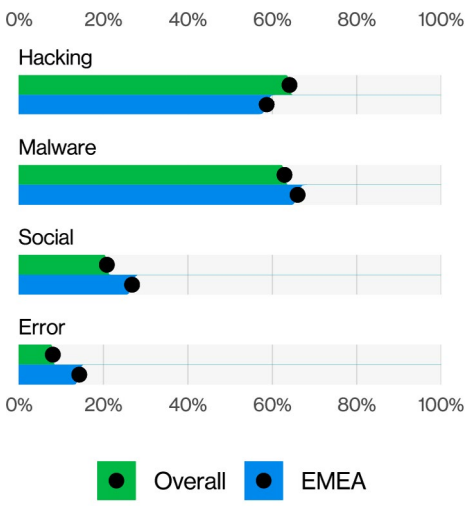

Figure 105. Top Action varieties in EMEA breaches (n for EMEA=6,060)

Error actions tell a slightly different story. In EMEA, Misdelivery leads the error category at 50%. While Misdelivery is also the top error type in the overall dataset (61%), it appears less frequently in EMEA by comparison. Loss, however, is notably higher in EMEA at 19%, versus 11% in the rest of the dataset.

Misconfiguration-related breaches come in at 17% for EMEA, compared to 13% overall. Of these three, Misconfiguration is arguably the most worrisome, even if it's less common than Misdelivery. Accidentally sending an attachment to the wrong recipient and exposing a small amount of data is unfortunate, but spinning up an unsecured cloud database holding terabytes of PII is the kind of mistake that can keep incident responders and legal teams up at night.

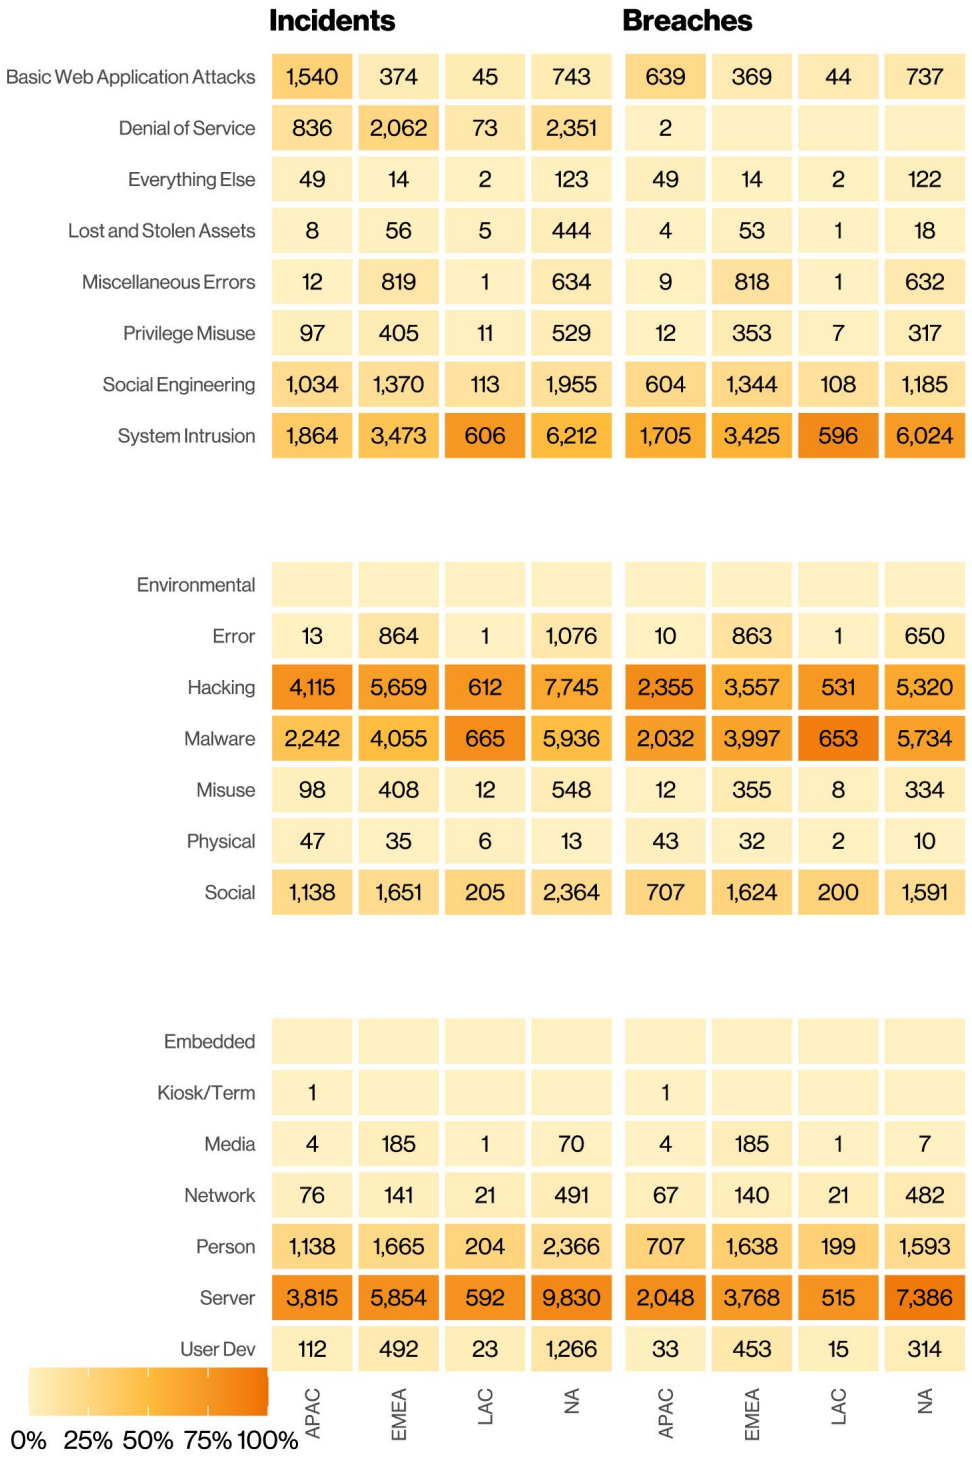

Figure 106. Incidents and breaches by region

Critical economic infrastructure

Any discussion regarding cybercrime in EMEA this year should probably include some mention of the massive 2025 attack on Jaguar Land Rover (JLR).[115] This incident represents the most economically damaging cyberattack in U.K. history, with an estimated £1.9 billion in loss due to a grueling five-week business interruption. This disruption rippled downstream to impact approximately 5,000 entities, demonstrating that incidents that affect an organization's supply chain can quickly escalate into loss for others involved, other than the criminals presumably.

The JLR breach was damaging to the economy as a whole. This was a contributing factor to the U.K. gross domestic product missing its projection by 0.1%,[116] which could be part of the reason why the government intervened with loans to support JLR and their supply chain. This makes us think about the real significance of critical infrastructure with today's deep economic interconnection and how to best support the private sector, in general, against attacks of this magnitude.

As one of their response measures, the U.K. National Cyber Security Centre (NCSC) has released a campaign to incentivize companies to increase their cyber resilience,[117] with a focus on recovery planning and execution after a damaging attack.

115. bbc.com/news/articles/cy9pdld4y81o
116. reuters.com/world/uk/uk-economy-grows-01-q3-2025-11-13
117. ncsc.gov.uk/campaigns/cyber-resilience

Wrap-up

/08

You are now free to return to your organization's ongoing crisis management activities, and thanks for reading.

Whether you used this report to sharpen your defensive strategy or simply dropped by to find fault, we sincerely hope you found it useful.

Each year, our goal is the same: to provide a compass that can assist you in navigating an increasingly chaotic environment by drawing attention to the threats that could be the most impactful to your organization and supporting your decision-making when deploying resources. We hope that, to some degree, we have accomplished that goal.

As we do every year, we owe a massive debt of gratitude to our data contributors. Without their willingness to share their data and expertise, this report simply wouldn't exist.

While we are handing out thank-you's, we thank you, our loyal readers, for taking the time to support us by reading and sharing this document and by asking questions or making suggestions that enable us to remain relevant in a world that refuses to ever stand still.[118]

Looking toward next year, we can't help but notice a milestone on the horizon: 2027 will mark the 20th anniversary of the DBIR! We aren't quite sure where the time (or our hair) went, but we are already hard at work ensuring that our platinum anniversary edition is the most impactful one yet.

On behalf of the entire team, we wish you a secure and prosperous year. Finally, as always, stay vigilant, stay out of the headlines and—most importantly—stay in touch.

118. Just a few weeks would be nice.

Year in review

Monthly snapshot as reported by the VTRAC Monthly Intelligence briefings and kindly provided by Steven Baskerville, Darrin Kimes and Jim Meehan from the VTRAC team

January	The identity and edge assault: The year opened with a strategic shift from traditional endpoints to core identity and edge infrastructure. Chinese state-sponsored actor Silk Typhoon compromised the U.S. Treasury Department by stealing a BeyondTrust security key, granting remote access to classified workstations. Simultaneously, the J-Magic campaign targeted Juniper routers, utilizing a custom cd00r backdoor that scanned for "magic packets" to establish reverse shells on VPN gateways. High-severity zero-days in Ivanti Connect Secure (CVE-2025-0282) were weaponized by UNC5337 to deploy the SPAWN malware ecosystem, providing persistent, unauthenticated access to hundreds of enterprise networks. This month also saw a massive phishing campaign hijack 35 Chrome extensions, injecting data-stealing code that impacted 2.6 million users.
February	Bypassing the identity perimeter: Adversaries focused on the "Identity Crisis," utilizing Device Code Authentication phishing to bypass MFA for Microsoft 365 accounts by impersonating high-level government officials. The Akira ransomware group demonstrated a novel persistence tactic by exploiting unsecured webcams to move laterally and encrypt VMware ESXi shares while remaining invisible to Windows-based endpoint detection and response systems. Law enforcement executed Operation PHOBOS AETOR, dismantling the 8Base/Phobos ransomware infrastructure in Thailand, though the group's affiliates rapidly pivoted to new Malware as a Service (MaaS) models. The month also saw the viral emergence of DeepSeek-R1, which was immediately met with 100% successful jailbreaking campaigns to generate malicious code.
March	Cascading supply chain failures: Supply chain fragility took center stage as a cascading breach of GitHub Actions exposed secrets for more than 23,000 repositories. China-nexus group UNC3886 (Weaver Ant) demonstrated extreme technical depth by maintaining network access for four years and bypassing Juniper's kernel-based file integrity (CVE-2025-21590). Law enforcement seized the Garantex crypto exchange after it processed $96 billion in illicit transactions. Meanwhile, the discovery of BADBOX 2.0 revealed a botnet of more than one million infected connected TVs used for systemic ad fraud. North Korea officially launched Research Center 227, a unit dedicated to developing AI-driven offensive hacking capabilities.
April	High-leverage extortion: Threat actors pivoted toward high-leverage sectors where downtime causes immediate systemic pressure. U.K. retailers Marks & Spencer, Co-op and Harrods were hit by ransomware—the Co-op breach alone compromised the personal data of 6.5 million members. A hacker known as "rose87168" claimed a massive breach of Oracle Cloud, exfiltrating 6 million records from 140,000 tenants. A critical zero-day in CrushFTP (CVE-2025-31161) allowed unauthenticated administrative takeovers via AWS4-HMAC race conditions.
May	The zero-day sprint: A sharp escalation in zero-day activity saw 15 such flaws added to the KEV catalog. The Russian Qilin ransomware group and China-nexus actors exploited a critical SAP NetWeaver unauthenticated file upload flaw (CVE-2025-31324) weeks before disclosure to deploy webshells globally. Russian hacktivists (NoName057(16)) launched persistent DDoS attacks against Dutch and Romanian state websites in retaliation for military aid to Ukraine. Ivanti Endpoint Manager Mobile (EPMM) software was hit by a zero-day exploit chain (CVE-2025-4427/4428) delivered via AWS S3 buckets to gain remote code execution on managed devices.
June	Blurring state and criminal lines: The distinction between financial extortion and state-sponsored data collection vanished. ShadowPad variants were deployed by China-aligned FamousSparrow against research institutes in the U.S. and Mexico. Operation Endgame disrupted Lumma Stealer and DanaBot infrastructure across 1,300 domains, though developers restored MaaS operations within days. A massive 631GB database leak exposed four billion user records from Chinese platforms WeChat and Alipay. French authorities arrested five operators of the BreachForums platform, including individuals associated with the ShinyHunters group.

July	Systemic infrastructure fragility: Fragility was underscored by the SharePoint ToolShell zero-day chain (CVE-2025-53770), allowing Chinese actors such as Linen Typhoon to gain unauthenticated access to hundreds of high-value networks. Citrix NetScaler faced a second crisis with CitrixBleed 2 (CVE-2025-5777), which was exploited as a zero-day to leak session tokens from critical infrastructure in the Netherlands. In Brazil, the Datzbro Android trojan began utilizing AI-generated Facebook events to trick elderly users into device-takeover attacks.
August	AI and infrastructure sabotage: AI moved from theoretical research into offensive implementation. MITRE revealed LameHug, an APT28 experiment using Alibaba's Qwen LLM to generate polymorphic malware code on demand. ShinyHunters (UNC6040) launched a massive supply-chain campaign, exploiting compromised Salesloft Drift OAuth tokens to pivot into the Salesforce instances of major firms such as Google, Zscaler and Cisco. The PromptLock malware emerged as the first AI-powered ransomware to generate cross-platform encryption scripts dynamically via local LLMs. Pro-Russian hackers were suspected of sabotaging a Norwegian dam, breaching control systems to release water.
September	The industrial impact: The industrial sector faced its costliest cyber event in U.K. history as a ransomware attack on Jaguar Land Rover (JLR) halted production for five weeks, causing £1.9 billion in damages. Amazon revealed it thwarted more than 1,800 North Korean "remote worker" infiltration attempts by identifying a unique 110 ms keystroke input lag. A self-replicating npm worm called Shai-Hulud compromised more than 500 packages to exfiltrate developer credentials and GitHub access tokens.
October	Volumetric and hypervisor warfare: Record-breaking volumetrics defined the month as the Aisuru botnet (300,000 IoT hosts) launched a record 29.7 Tbps DDoS attack, nearly doubling previous peaks. State-sponsored actors exploited Cisco ASA zero-days (CVE-2025-20333) to deploy the LINE VIPER and RayInitiator malware families. Federal authorities seized nearly 130,000 Bitcoins (approx. $15 billion) from the Cambodian Prince Group, targeting a massive investment fraud and human trafficking empire. Analysis found that 29% of KEV vulnerabilities were attacked before public disclosure this year.
November	The vishing and SaaS siege: Vishing (voice phishing) surged, with the FakeCall malware intercepting calls on infected mobile devices to steal banking credentials. ShinyHunters expanded their SaaS siege, breaching Gainsight to access 285 additional Salesforce instances. A nation-state actor gained long-term access to F5's development environment, exfiltrating BIG-IP source code and undisclosed vulnerability data. Airstalk malware emerged in a supply chain attack that misused MDM APIs as a "dead drop" for C2 communication.
December	Mass exploitation of modern stacks: The year closed with the widespread exploitation of React2Shell (CVE-2025-55182), an RCE vulnerability in React Server Components exploited by China-nexus groups to deploy backdoors across 39% of cloud environments. North Korean hackers reached a record annual theft of $2.02 billion in cryptocurrency. The discovery of VoidLink, a malware framework written in six days by an AI agent, marked a point of no return for automated threat development.

Appendices

/09

Appendix A: Methodology

One of the things readers value most about this report is the level of rigor and integrity we employ when collecting, analyzing and presenting data. Knowing our readership cares about such things and consumes this information with a keen eye helps keep us honest. Detailing our methods is an important part of that honesty.

To begin with, we would like to remind our readers that science comes in two flavors: creative exploration and causal hypothesis testing. The DBIR is squarely in the former. While we may not be perfect, we believe we provide the best obtainable version of the truth based on the datasets we have available (to a given level of confidence and under the influence of biases acknowledged later). However, proving causality is best left to randomized control trials. The best we can do is correlation. And while correlation is not causation, they are often related to some extent, and often useful.

Non-committal disclaimer

We would like to reiterate that we make no claim that the findings of this report are representative of all data breaches in all organizations at all times. Even though we believe the combined records from all our contributors more closely reflect reality than any of them in isolation, it is still a sample. And although we believe many of the findings presented in this report to be appropriate for generalization (and our conviction in this grows as we gather more data and compare it to that of others), bias still exists.

The DBIR process

Our overall process remains intact and largely unchanged from previous years.[119]

All incidents included in this report were reviewed and converted (if necessary) into the VERIS framework to create a common, anonymous aggregate dataset. If you are unfamiliar with the VERIS framework, it is short for Vocabulary for Event Recording and Incident Sharing, it is free to use, and links to VERIS resources can be found throughout this report.

The collection method and conversion techniques differed among contributors. In general, three basic methods (expounded below) were used to accomplish this:

1. Direct recording of paid external forensic investigations and related intelligence operations conducted by Verizon using the VERIS Webapp
2. Direct recording by partners using VERIS
3. Converting partners' existing schema into VERIS

All contributors received instruction to omit any information that might identify organizations or individuals involved.

Some source spreadsheets are converted to our standard spreadsheet formatted through automated mapping to ensure consistent conversion. Reviewed spreadsheets and VERIS Webapp JSON are ingested by an automated workflow that converts the incidents and breaches within into the VERIS JSON format as necessary, adds missing enumerations, and then validates the record against business logic and the VERIS schema. The automated workflow subsets the data and analyzes the results. Based on the results of this exploratory analysis, the validation logs from the workflow and discussions with the partners providing the data, the data is cleaned and reanalyzed. This process runs nightly for roughly two months as data is collected and analyzed.

Incident data

Our data is non-exclusively multinomial, meaning that a single feature, such as "Action," can have multiple values (e.g., "Social," "Malware" and "Hacking"). This means that percentages do not necessarily add up to 100%. For example, if there are five botnet breaches, the sample size is five. However, because each botnet used phishing, installed keyloggers and used stolen credentials, there would be five Social actions, five Hacking actions and five Malware actions, adding up to 300%. This is normal, expected and handled correctly in our analysis and tooling.

Another important point is that when looking at the findings, "unknown" is equivalent to "unmeasured." Which is to say that if a record (or collection of records) contains elements that have been marked as "unknown" (whether it is something as basic as the number of records involved in the incident or as complex as what specific capabilities a piece of malware contained), it means that we cannot make statements about that particular element as it stands in the record—we cannot measure where we have too little information. Because they are unmeasured, they are not counted in sample sizes. The enumeration "Other," however, is counted because it means that the value was known but not part of VERIS (or not one of the other bars if found in a bar chart). Finally, "Not Applicable" (normally "n/a") may be counted or not counted depending on the claim being analyzed.

We make liberal use of confidence intervals to allow us to analyze smaller sample sizes. We have adopted a few rules to help minimize bias in reading such data. Here we define "small sample" as fewer than 30 samples.

1. Sample sizes smaller than five are too small to analyze.

119. As does this sentence

2. We won't talk about count or percentage for small samples. This goes for figures, too, and is why some figures lack the dot for the median frequency.
3. For small samples, we may talk about the value being in some range or values being greater/less than each other. These all follow the confidence interval approaches listed previously.

Incident eligibility

For a potential entry to be eligible for the incident/breach corpus, a few requirements must be met. The entry must be a confirmed security incident defined as a loss of confidentiality, integrity or availability. In addition to meeting the baseline definition of "security incident," the entry is assessed for quality.

We create a subset of incidents that pass our quality filter. The details of what is a "quality" incident are:

- The incident must have at least seven enumerations (e.g., threat actor variety, threat action category, variety of integrity loss) across 34 fields OR be a DDoS attack. Exceptions are given to confirmed data breaches with fewer than seven enumerations.
- The incident must have at least one known VERIS threat action category (e.g., Hacking, Malware).

In addition to having the level of details necessary to pass the quality filter, the incident must be within the timeframe of analysis (Nov 1, 2024, to Oct 31, 2025, for this report). The 2025 caseload is the primary analytical focus of the report, but the entire range of data is referenced throughout, notably in trending graphs. We also exclude incidents and breaches affecting individuals that cannot be tied to an organizational attribute loss. If your friend's laptop was hit with ransomware while downloading a game cheat, it would not be included in this report.

Lastly, for something to be eligible for inclusion into the DBIR, we have to know about it, which brings us to several potential biases.

Acknowledgement and analysis of bias

Many breaches go unreported (though our sample does contain some of those, as well). Many more are as yet unknown by the victim (and thereby unknown to us). Therefore, until we (or someone) can conduct an exhaustive census of every breach that happens in the entire world each year (our study population), we must use sampling. Unfortunately, this process introduces bias.

The first type of bias is random bias introduced by sampling. This year, our maximum confidence is +/- 0.7% for incidents and +/- 0.9% for breaches, which is related to our sample size. Any subset with a smaller sample size is going to have a wider confidence margin. We've expressed this confidence in the complementary cumulative density (slanted) bar charts, hypothetical outcome plot (spaghetti) line charts and quantile dot plots. However, sometimes the nature of non-incident data we may be working with is not conducive to this confidence level analysis, and we might have some plain vanilla bar and line charts throughout the report. More on non-incident data in the next section.

The second source of bias is sampling bias. We strive for "the best obtainable version of the truth" by collecting breaches from a wide variety of contributors. Still, it is clear that we conduct biased sampling. For instance, some breaches, such as those publicly disclosed, are more likely to enter our corpus, while others, such as classified breaches, are less likely.

We also acknowledge that some types of breaches that are very common in a specific analysis period – looking at you, Ransomware – might end up being overrepresented due to the vast availability of samples. We often try to point it out in the report when that is the case.

The third source of bias is confirmation bias. Because we use our entire dataset for exploratory analysis, we cannot test specific hypotheses. Until we develop a collection method for data breaches beyond a sample of convenience, this is probably the best that can be done.

As stated earlier, we attempt to mitigate these biases by collecting data from diverse contributors. We follow a consistent multiple-review process and when we hear hooves, we think horses, not zebras.[120] We also try and review findings with subject matter experts in the specific areas ahead of release.

Non-incident data

Since the 2015 issue, the DBIR has included data that requires analysis that does not fit into our usual categories of "incident" or "breach." Examples of non-incident data include malware, vulnerability management, phishing, DDoS, internet-wide honeypots, internet-wide scanning and other types of data. The sample sizes for non-incident data tend to be much larger than the incident data but from fewer sources. We make every effort to normalize the data (for example, weighting records by the number contributed from the organization so all organizations are represented equally). We also attempt to combine multiple partners with similar data to conduct the analysis wherever possible. Once analysis is complete, we try to discuss our findings with the relevant partner or partners so as to validate the findings against their knowledge of the data and make sure we are representing it correctly.

120. A unique finding is more likely to be something mundane, such as a data collection issue, than an unexpected result.

Appendix B: U.S. Secret Service

By Assistant Special Agent in Charge Richard Hersh III, Digital Forensics Incident Response-Network Intrusion PM Bernard Wilson, and Management and Program Analyst Stephen Hampton, United States Secret Service

The autonomous adversary

The digital battlefield continues to evolve. Agentic AI (autonomous AI capable of independent action) is redefining cybercrime by creating adversaries that can operate without human limits. Traditionally, cybercriminals relied on human effort and technical skill to execute attacks. Now, Agentic AI systems can automate every stage of cybercrime: reconnaissance, phishing, data theft, and even laundering stolen and illicit assets. The United States Secret Service is on the front lines with an unwavering commitment to counter cyber threats.

Agentic AI has the potential to lower the barrier for sophisticated cyberattacks while expanding their scope. Agentic AI can run persistent campaigns, adapt tactics in real time, and target thousands, if not millions, of victims simultaneously. As a result, cyberattacks are often faster, smarter, and more relentless. Agentic AI doesn't just scale up old tricks, it invents new ones, automates deception, and exploits vulnerabilities at a pace that overwhelms traditional defenses.

This shift is both profound and urgent for law enforcement, and a call to innovate with exceptional operational technologies. Adversaries that never sleep can test law enforcement capabilities and incident response expertise. This new type of cybercriminal operates without human fatigue or skill limitations. Agentic AI can generate convincing scam messages, impersonate trusted contacts, and orchestrate complex attacks seamlessly.

These threats are increasing in volume and diversity, and even unskilled criminals can launch sophisticated campaigns with just a few queries to an AI platform. Law enforcement responders and investigators face a clear challenge: they must adapt or risk falling behind.

Autonomous adversaries are pushing cybercrime into a new era, one where attacks are limited only by the imagination of the algorithms behind them. The Secret Service has a mandate to investigate these cybercrimes under the Computer Fraud and Abuse Act, codified at Title 18, United States Code, Section 1030.[121] Under this authority, the Secret Service continues to evolve investigative techniques, advance analytics, and leverage collaborative expertise across public and private sectors to counter AI-driven threats. The Secret Service is leaning into this challenge, ensuring that even as adversaries become more autonomous, law enforcement remains agile, innovative, and relentless in pursuit of cybercriminals.

As a defense tool, Agentic AI can deploy real-time monitoring, automate threat detection, and implement innovative response platforms, which can significantly reduce Mean Time to Respond/Repair (MTTR) and are essential tools in the fight against autonomous adversaries.

The stakes have never been higher. The battle is just beginning. The Secret Service refuses to stand still.

121. uscode.house.gov/view.xhtml?req=(title:18%20section:1030%20edition:prelim)

Appendix C: Using the DBIR for Security Risk Decisions

By Tony Martin-Vegue, Cyber Risk Expert

Every year, thousands of security professionals follow the same ritual: download the DBIR, read the findings, and use them to understand trends, track the evolving threat landscape, inform security decisions, and support planning.

The Verizon Data Breach Investigations Report is one of a kind, a longitudinal study of security trends that has been ongoing for 19 years. It's a gold mine of useful data for organizations, large and small. However, some security professionals struggle to get from "Phishing was involved in 14% of healthcare breaches" to actionable decisions. How do you use evidence like this to improve your security and risk decisions?

This appendix provides practical advice on using DBIR data in a way that withstands scrutiny within your organization. You'll understand where the data comes from, what it means, and how to use it.

Understanding the statistics

One of the first things to understand when reading the report is how to interpret the statistics. For example, when the DBIR reports that "ransomware was present in 48% of breaches," that does not mean your organization or your sector has a 48% chance of being hit by ransomware.

Mistaking the statistic for probability is a common mistake, and it can lead to misguided risk decisions.

Unpacking what the statistic means is the key. Think of it this way: when we say ransomware was present in 48% of breaches, what we're really saying is "among organizations that got breached AND detected it AND reported it, or had it reported by someone else, 48% of those breaches involved ransomware."

Think for a moment about the filter that creates. For something to appear in the DBIR dataset, the attack had to succeed, someone had to notice it, it had to meet eligibility/quality criteria, and it had to be available to contributors.

Each one of those steps is important. Each one creates a filter for what you see. So, it's not a probability of a data breach. It's a statistic about observed organizations that experienced an incident serious enough to end up in the dataset, and among those, 48% involved ransomware.

The DBIR reports detected and reported data breaches and the types and patterns of those incidents. It does not capture attacks that failed, activity that was blocked or disrupted before causing harm, or incidents that went undetected or never met reporting criteria. All of those nuances matter for understanding both your risk and the statistics the DBIR is reporting.

If you miss this distinction, everything else you interpret about the report will be off.

Using the DBIR as your baseline

The DBIR splits things out by industry, size, and region for a reason: threat landscapes look different depending on where you sit. What's common for a hospital might be rare for a retailer. What hits a small business isn't the same as what goes after an enterprise. Your first step is to find your relevant demographic in the report: your industry vertical, your organization size, and your geographic region.

With tens of thousands of incidents analyzed over 19 years, the DBIR provides a baseline for what has historically happened to organizations in your category. This provides very useful context and tells you what an attack would look like at your organization.

Understanding the baseline and its limitations

There's an important caveat here: the DBIR only captures organizations that got breached, detected it, and had it reported. This creates selection bias, which means the sample (breached organizations in the report) may be unrepresentative of the whole population (all organizations). By definition, these organizations in the report had defenses that failed. Does that mean they all had worse security than average? Not necessarily. The sample includes organizations across the spectrum: some with weak security that were easy targets, some with decent security that got unlucky, and some with strong security that faced sophisticated, determined attackers.

The very large sample size across many sectors provides more comprehensive coverage. You're seeing a reasonably wide distribution of outcomes in the real world, not just the worst-case scenarios. However, you're not seeing organizations that successfully defended against attacks, and you're not seeing breaches that went undetected.

What this means for you: the DBIR shows you patterns in breaches when they occur in your sector. It tells you what attack methods succeeded, what vulnerabilities got exploited, and what controls failed. It doesn't tell you your overall probability of being breached. That depends on your specific situation.

Now the question becomes: when you look at these patterns, are you better prepared than what's typical in these incidents, worse, or about the same? Do you have stronger controls than those that failed in these breaches? Weaker detection capabilities? Different threat exposure?

Understanding where you stand relative to these patterns is how you turn DBIR statistics into useful information for your organization. Practical ways to assess where you stand include conducting internal control assessments and comparing your security posture against sector benchmarks. Many research firms, vendors, and industry reports publish control benchmarking data for different sectors, and these can help you understand whether your controls are stronger, weaker, or similar to what's typical in organizations like yours. This can help you adjust the baseline DBIR statistic up or down based on your findings.

What evidence helps you adjust the baseline

Once you've identified your relevant DBIR baseline, the next step is figuring out whether you should adjust it up or down for your organization. Here's what evidence matters:

Strong evidence:

Your own history. This is the strongest evidence you have. If you've been successfully phished multiple times, that's direct evidence of your vulnerability to phishing attacks. If you've been hit by ransomware before, that gives you concrete evidence of your exposure. Your track record provides real data about how your defenses perform under actual attack conditions.

What attackers find when they test you. Pen test results, red team findings, and bug bounty reports show you what external testers discover about your defenses. If they're finding critical vulnerabilities, that's evidence that you may be more vulnerable than the baseline. If they're coming up mostly empty, that suggests stronger defenses.

What controls do you have versus what failed in breaches? DBIR findings often describe what was missing or what broke in successful attacks: No MFA, unpatched critical systems, weak passwords, successful social engineering, etc. Compare that list to your environment, and if you have some of the same gaps, you may be more vulnerable than the baseline. If you've addressed those common failure points, you may be better positioned.

Your ability to detect problems. The DBIR only captures breaches that were detected. If your detection capabilities are limited, you might have exposure you don't know about yet. Strong detection and monitoring capabilities suggest you're better positioned than organizations that miss breaches entirely.

Supporting evidence:

How you stack up against baseline requirements. If you're barely meeting compliance minimums, you're likely close to the baseline. If you're exceeding requirements significantly, that suggests stronger controls.

What your peers tell you. Conversations with similar organizations can help you understand whether your security posture is typical, ahead, or behind what's common in your sector.

What insurers care about. Cyber insurance questionnaires often reflect what insurers see as meaningful risk factors based on their claims data. Where you stand on those factors provides useful context.

This does not need to be an overly onerous and time-consuming exercise. Gather just enough evidence to make a reasonable judgment about whether you should adjust the baseline DBIR statistics up or down for your specific situation.

Applying this in practice

A hospital reads the DBIR and sees that credential compromise following successful phishing is common in healthcare breaches. They ask themselves: Should we adjust the baseline up or down for our organization?

They look at the evidence:

- MFA is deployed on email and the EHR system, but not on VPN
- Successfully phished twice in the past eighteen months
- Last pen test showed weak password practices and a 28% click rate on simulated phishing
- Security awareness training happens once a year
- Peer hospitals that got breached had similar security programs

The evidence suggests they're more vulnerable than the baseline. That's never a fun conversation with leadership, but it changes decisions.

In our example, phishing becomes the top priority. They extend MFA to VPN access, move to continuous security awareness training, and implement monthly phishing simulations.

Six months later, simulated phishing click rates dropped from 28% to 8%. The next pen test shows measurable improvement. They've improved their position relative to the baseline, and their risk profile has changed because the security projects have moved them to a stronger position than the industry baseline.

Making it actionable

The DBIR gives you 19 years of data about what happens when organizations in your sector get breached. Use it as your starting point, then gather evidence about your specific situation to determine if your risk is higher or lower than that baseline.

When you read about patterns in the data, ask yourself: are we better prepared than this baseline, worse, or about the same? If you're worse, you know what needs to change. If you're better, you should be able to point to specific evidence showing why.

That's how you turn industry statistics into actionable risk decisions for your organization.

Appendix D: Contributing organizations

A

Abstract

Akamai Technologies

Ankura

Anthropic

Apura Cyber Intelligence

Archer Hall

Atos

B

bit-x-bit

Bitsight

Brand Defense

Breachlock

Bridewell

C

Censys, Inc.

Center for Internet Security (CIS)

Cequence Security

CERT – European Union (CERT-EU)

Check Point Software Technologies Ltd.

Coalition

Compass Security

Coveware

COWBELL

Cyber Security Agency of Singapore

Cybersecurity Infrastructure Security Agency (CISA)

CyberSecurity Malaysia, an agency under the Ministry of Communications and Multimedia (KKMM)

Cybersixgill

D

DarkWeb IQ

Defense Counterintelligence and Security Agency (DCSA)

Department of Home Affairs

Digital.ai

DigitalMint

DomainTools

Dragos, Inc

DTEX

E

Edgescan

Emergence Insurance

Empirical Security

Enduir

Enzoic

EUROCONTROL

F

Fastly

Federal Bureau of Investigation – Internet Crime Complaint Center (FBI IC3)

F-Secure

Flare

Flashpoint

G

Global Resilience Federation

GreyNoise

H

HackNotice

Halcyon

Hoxhunt

I

ImmuniWeb

Infoblox

Information Commissioner's Office (ICO)

Irish Reporting and Information Security Service (IRISS-CERT)

J

JPCERT/CC

K

K–12 Security Information Exchange (K–12 SIX)

Keep Aware

Keepnet Labs

KnowBe4

KordaMentha

L

LayerX

Legal Services Information Sharing and Analysis Organization (LS-ISAO)

Lookout

M

Maritime Transportation System ISAC (MTS-ISAC)

Mimecast

N

National Crime Agency

National Cyber-Forensics & Training Alliance (NCFTA)

Netclean

NetDiligence®

NETSCOUT

O

Office of the National Cyber Security Agency, Thailand

Okta

Onapsis

OpenText Cybersecurity

P

Proofpoint

Q

Qualys

R

Recorded Future, Inc.

RedHunt Labs

ReversingLabs

S

SAFE

Security Scorecard

Shadowserver Foundation

Shodan

Sistemas Aplicativos

Six Degrees

Sophos

SpecterOps

Swisscom

T

Tenable

Tenchi Security

Thales

The CWE Program

The DFIR Report

Tidal Cyber

Triskele Labs

TRM Labs

U

U.S. Secret Service

V

Veracode

VERIS Community Database

Verizon Cyber Security Consulting

Verizon DDoS Defense

Verizon Network Operations and Engineering

Verizon Threat Research Advisory Center (VTRAC)

Verizon VTRAC Labs

W

Wabtec

Wiz

Z

Zscaler


abstract
Akamai
ankura
ANTHROPIC
apura
ARCHERHALL
AIM HIGH
Atos
bit-x-bit
BITSIGHT
BRANDEFENSE
breachlock
Bridewell
censys
ARC
CIS
Center for Internet Security
CEQUENCE
CERT-EU
CHECK POINT
Coalition
COMPASS SECURITY
COVEWARE
cowbell
CSA SINGAPORE
Cyber Security Agency of Singapore
CISA
CyberSecurity MALAYSIA
cybersixgill
A BITSIGHT Company
DARKWEB IQ
DEFENSE COUNTERINTELLIGENCE AND SECURITY AGENCY
Australian Government
Department of Home Affairs
digital.ai
digitalmint
cyber
DomainTools
DRAGOS
DTEX
edgescan
emergence
EMPIRICAL
ENDUIR
ENZOIC
EUROCONTROL
EATM-CERT
European Air Traffic Management
Computer Emergency Response Team
fastly
IC3
F-Secure
flare
FLASHPOINT
GLOBAL RESILIENCE FEDERATION
GREYNOISE
HackNotice
halcyon
HOXHUNT
ImmuniWeb
AI for Application Security
infoblox
ico.
Information Commissioner's Office
IRISS
Irish Reporting and Information Security Service
JPCERT CC

K12 SIX	keep aware	keepnet	knowbe4	KordaMentha	LayerX
LEGAL SERVICES ISAO	Lookout	Maritime Transportation System Information Sharing and Analysis Center	mimecast	NCA National Crime Agency	NCFTA
NetClean.	NetDiligence	NETSCOUT	NCSA สกมช	okta	ONAPSIS
opentext™	proofpoint.	Qualys.	Recorded Future®	REDHUNT LABS DISCOVER. ATTACK. REPEAT.	REVERSINGLABS
SAFE	SecurityScorecard	SHADOWSERVER	SHODAN	SISAP Sistemas Aplicativos	Six Degrees
SOPHOS	SPECTEROPS Creators of BloodHound	swisscom	tenable	TENCHI	THALES Building a future we can all trust
CWE	THE DFIR REPORT	TIDAL CYBER	Triskele Labs	TRM	
VERACODE		Verizon Cyber Security Consulting	Verizon DDoS Defense	Verizon Network Operations and Engineering	Verizon Threat Research Advisory Center (VTRAC)
Verizon VTRAC Labs	Wabtec CORPORATION	WIZ	zscaler™