

2026

Trend Report

FOR THE CYBERSECURITY LANDSCAPE



HUNT &
HACKETT

“Changes in the threat landscape show that focus needs to shift from the perimeter to identities and digital supply chains.”

Ronald Prins,
Co-founder at Hunt & Hackett

Table of Contents

Foreword	1
Chapter 1: Our view from the frontlines	4
1.1 Key Trends	4
1.2 A snapshot of our SOC	5
1.3 Insights from Incident Response	13
1.4 EDR variability in practice	15
Chapter 2: The evolving attacker playbook	16
2.1 Key Trends	16
2.2 Paying the interest on technical debt	17
2.3 The expanding attack surface	18
2.4 Identity as a leading attack vector	25
2.5 Gen-AI exacerbates asymmetry between attack and defense	29
Chapter 3: Cyber at the geopolitical frontlines	33
3.1 Key Trends	33
3.2 The rise of hacktivism	34
3.3 Blurred lines	36
3.4 The Big Four	38
3.5 We're not seeing everything	44
Chapter 4: The big picture	46
4.1 Key Trends	46
4.2 Digital sovereignty	47
4.3 The cybersecurity market	53
4.4 So, what should I do on Monday?	55
References	56
Appendix	61

Copyright © Hunt & Hackett BV All rights reserved. Nothing in this publication or on this internet website may be reproduced, stored in a computer database, in automatic and/or digital files, published, in any form or in any way, either electronically, mechanically, by means of photocopy, pictures, tapes or in any other way, without preceding explicit written permission of Hunt & Hackett BV.

Trademark Hunt & Hackett and the logo of Hunt & Hackett are trademarks of Hunt & Hackett BV. All other in this document published trademarks are owned by the corresponding named organisations.

Foreword

Welcome to Hunt & Hackett's 2026 Trend Report.

This second annual report aims to provide a clear, comprehensive view of today's threat landscape, so defenders know where to focus their efforts. That perspective is grounded in what our analysts see every day while monitoring customer environments and responding to incidents, combined with insights from experts across multiple disciplines. The report is underpinned by data from more than 54,000 security investigations carried out by our Security Operations Centre (SOC) and Incident Response team in 2025.

But understanding threats is only half the story. Geopolitical tensions have sharpened the debate about Europe's reliance on non-European cloud infrastructure and security platforms. As the pendulum swung toward more complex, interconnected, and outsourced IT environments, structural dynamics in the cybersecurity market shaped outcomes. These dynamics, and the interdependencies they created, may drive unwelcome results in the years ahead. That's why we also look in the mirror and examine the market forces that can undercut meaningful security improvement. This perspective rarely appears in industry reports, but it matters. Product-centric thinking, passive monitoring, and pricing models that force organizations to choose between data completeness and cost shape security outcomes as much as any attack technique.

Our goal is to offer a nuanced, practical view of what's really happening, both on the ground and in the market, so you can focus on the actions that matter most to you.

We hope you find the report useful. Thanks for reading.

Jurjen Harskamp
Co-founder and CEO

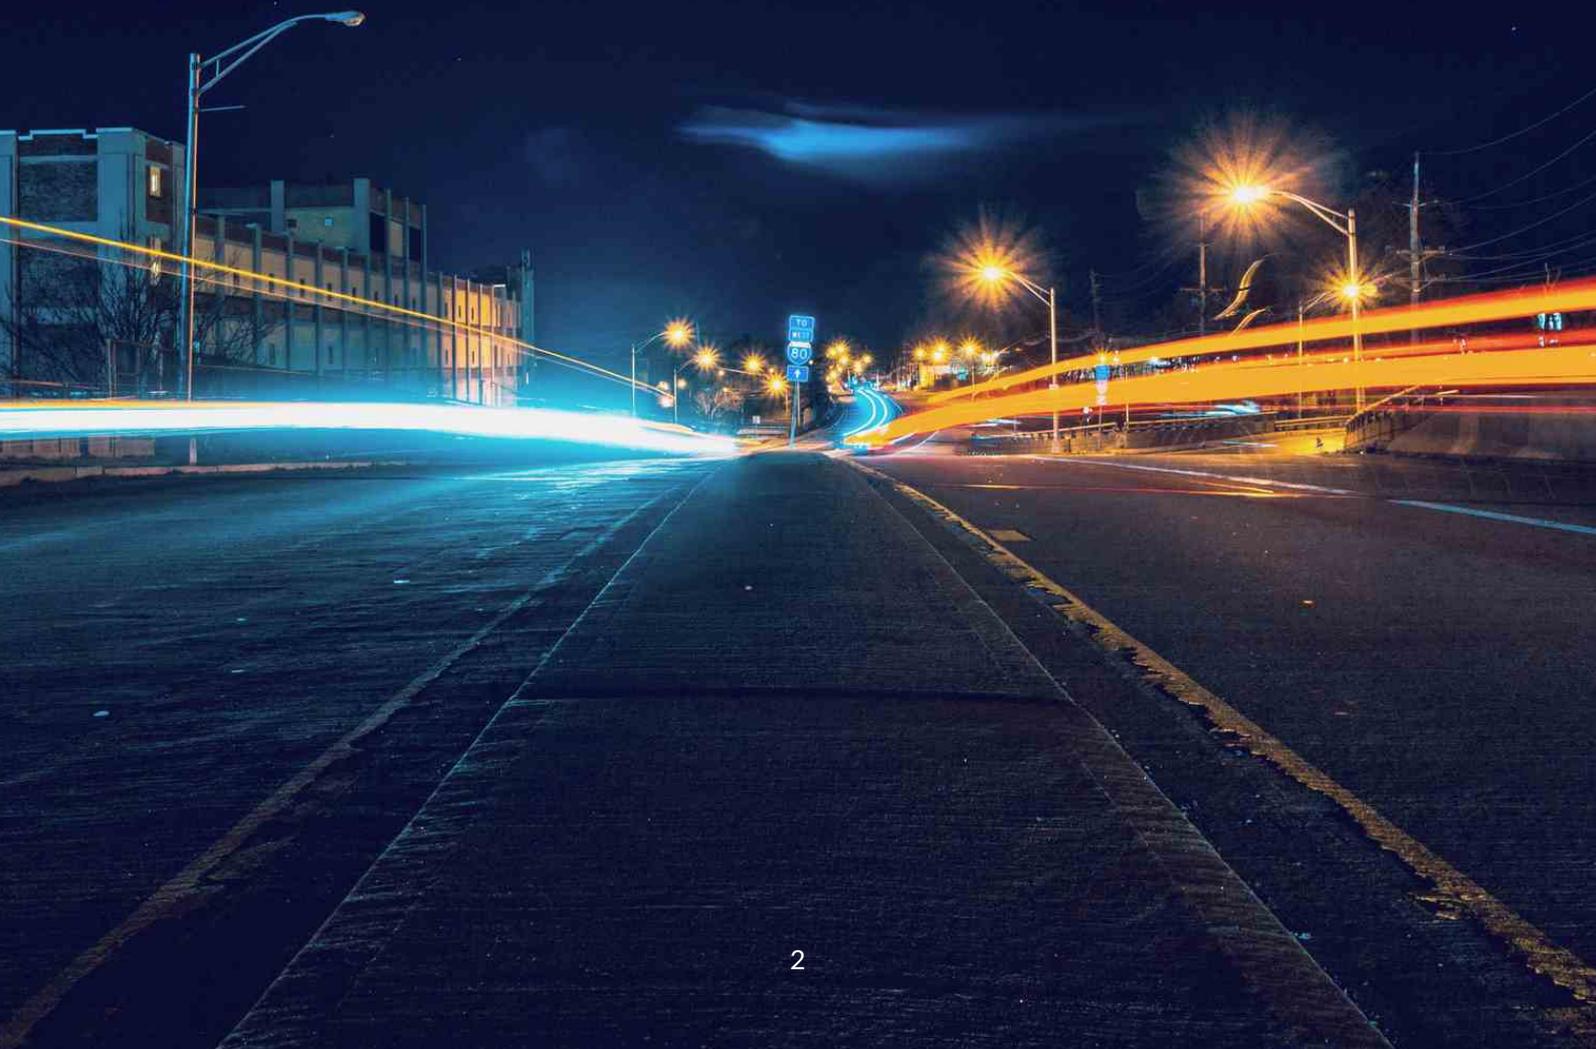


Ronald Prins
Co-founder



In 2026, the cybersecurity landscape is shaped by five key trends:

- **Financially motivated attacks remain a primary concern for Dutch organizations:** Ransomware was the most common incident type encountered by Hunt & Hackett's Incident Response team in 2025, followed by Business Email Compromise (BEC).
- **The attacker playbook is evolving incrementally:** Adversaries are increasingly exploiting technical debt, finding gaps in an ever-expanding attack surface, pivoting to identity-centric methods and employing generative AI to scale and accelerate their operations.
- **The boundaries between state actors, cybercriminals, and hacktivists continue to blur:** Groups increasingly share tools, masquerade as one another and conduct dual-motive operations, complicating attribution.
- **Geopolitical tensions have exposed Europe's deep dependence on US cloud infrastructure and security platforms:** Amid intensifying calls for digital sovereignty, experts shift the focus towards control points rather than absolute independence.
- **The cybersecurity market faces structural issues that undercut resilience:** Product-centricity, passive MDR, (semi-)closed systems, and misaligned SIEM economics make it difficult to achieve meaningful security improvements.



Thank you to our contributors



Erik Jonkman
Partner & Breach Counsel
at Kennedy Van der Laan

"Exploitation of edge devices, particularly vulnerable VPNs and firewalls, was a common theme in 2025."



Martijn Grooten
Threat Intelligence Analyst
and author of [Lapsed Ordinary](#) blog

"It's common to see threat actors copying from each other. ClickFix is just one example from the past year."



Stjepan Picek
Associate Professor at
Radboud University

"AI has shifted the cost-balance ratio for attackers, meaning 'low-value' targets have become more attractive."



Ellen Mok
Founder at Digitale
Doetank

"If you truly understand your stack and your dependencies, you can make conscious choices."



Daan Weggemans
Assistant Professor at
Leiden University

"Attackers are getting more professional, and their activities are becoming harder to recognize."



Rebecca Lumley
Strategic Threat
Intelligence Analyst at
Hunt & Hackett

"Abusing legitimate identities provides a stealthy way in for hackers - why break a window when you can just use a key?"



Jurjen Harskamp
Co-founder and CEO at
Hunt & Hackett

"Attackers exploit the gaps and delays created by technical debt that can't be resolved without operational risk."



Ronald Prins
Co-founder at Hunt &
Hackett

"Organizations install an EDR and think it will catch everything, but our data shows that this is not the case."



CHAPTER ONE

Our view from the frontlines

As a Managed Detection and Response (MDR) provider, we spend every day monitoring our customers' environments and responding to security incidents. This work, spanning both our Security Operations Center (SOC) and Incident Response (IR) engagements, generates a lot of data – data that reveals, at a granular level, how cyberattacks unfold in the real world.

While industry reports often rely on broad statistics, this chapter offers a bottom-up perspective grounded in what our analysts see every day. We believe this provides useful insights into not only the threats Dutch organisations are facing, but also the real-world challenges of defending against them.

Key trends

- **Financial gain was the primary motivation:** 71% of incidents handled by Hunt & Hackett's Incident Response (IR) team in 2025 were financially motivated. Ransomware was the leading attack type, accounting for 43% of IR engagements, followed by Business Email Compromise (BEC) at 29%.
- **Identity abuse was a core feature:** Credential theft and Active Directory (AD) exploitation featured prominently. Attackers employed Kerberoasting, AD database compromise, exploitation of weak passwords and Adversary-in-the-Middle phishing to bypass Multi-Factor Authentication (MFA).
- **The fundamentals were missing:** In most incident response engagements, the techniques used were widely known and detectable. What gave attackers room to operate was incomplete telemetry, lack of active monitoring and absence of fundamental security hygiene.

A snapshot of our SOC

In 2025, our Security Operations Centre (SOC) conducted more than 54,400 security investigations, with detections ranging from benign anomalies to genuinely malicious events. As this data is derived entirely from our customer base, it provides a direct view into the threats currently impacting Dutch organisations. To understand the attacker behaviours driving these investigations, we mapped them to the MITRE ATT&CK framework.

The most prevalent techniques, which are shown in the table on page eight, highlight a few clear trends. First, attackers typically gained their initial foothold by either targeting the human layer of the organization (User Execution, T1204) or by exploiting unpatched, internet-facing systems (Exploit Public-Facing Application, T1190). Once inside, attackers displayed a preference for living-off-the-land (LotL) techniques, abusing an organization's in-built system tools (Command and Scripting Interpreter, T1059) or hijacking legitimate signed binaries to run malicious code (System Binary Proxy Execution, T1218). To further their objectives, they sometimes brought in additional malicious tools (Ingress Tool Transfer, T1105) or used legitimate tools to blend in with normal administrative activity (Remote Access Tools, T1219).

This was typically followed by attempts to escalate privileges and evade defenses. Attackers achieved persistence by configuring malware to run automatically at startup (Boot or Logon Autostart Execution, T1547) or in response to system events (Event Triggered Execution, T1546). To move freely across the network, they stole credentials directly from the operating system (OS Credential Dumping, T1003) and actively disabled security controls to avoid detection (Impair Defenses, T1562).

T1204

User Execution was the most frequently observed technique (28%)

Methodology

The MITRE ATT&CK analysis presented in this section is based on more than 20,000 security investigations triggered by Hunt & Hackett's proprietary detection ruleset, constituting 38% of all investigations in 2025. Investigations triggered by third-party tools or customer-specific use cases are excluded from this analysis. From page nine onwards, findings are based on the full dataset comprising 54,400 investigations.

The resulting distribution of techniques is shaped by two primary factors: how often a technique is used in the wild, and whether we have rules in place to detect it. This can create a feedback loop, whereby we build more detections for techniques we encounter frequently, in turn increasing the likelihood of detecting them again. To avoid overfitting to our own dataset and maintain balanced coverage, we continuously track and study attacker methods across the wider threat landscape and adjust our rules accordingly.

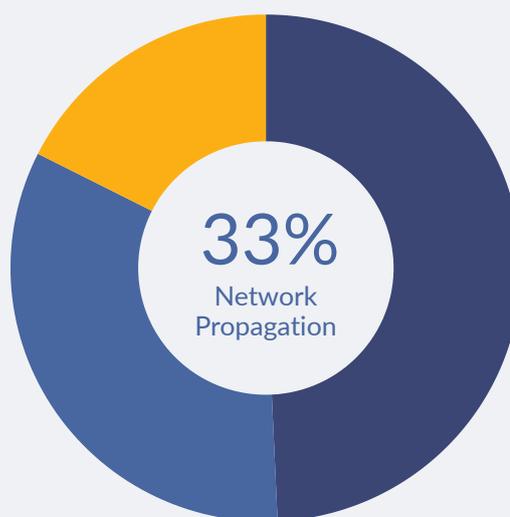
Assume breach requires visibility across the full attack chain

A mature security posture starts with an “assume breach” mindset: assume an attacker will eventually get a foothold, and focus on the ability to detect, investigate, and contain them as they move through the environment. That requires consistent visibility, not only at the perimeter but also across identity, endpoints, network traffic, cloud services, and critical business systems. This allows defenders to follow the attacker’s path from initial access to impact.

To see how our security investigations were distributed, we mapped more than 20,000 detections to MITRE ATT&CK tactics and categorized the resulting distribution into three phases: “In” (Initial Foothold); “Through” (Network Propagation); and “Out” (Action on Objectives). The concentration of detections is shown in the heatmap on the following page.

Heatmap interpretation

The heatmap on the next page depicts how detected security events relate to two widely recognized frameworks in our industry: the Unified Kill Chain and MITRE ATT&CK. We have chosen to present the data in this way to show not only the distribution of techniques, but also how they relate to the broader attack chain. Please note that it is not a strict one-to-one mapping. In cases where precision and clarity were at odds, we have prioritized readability, recognizing that data mapping and attack classification are areas the entire cybersecurity industry continues to refine.



A third of detections occurred in the “Through” phase

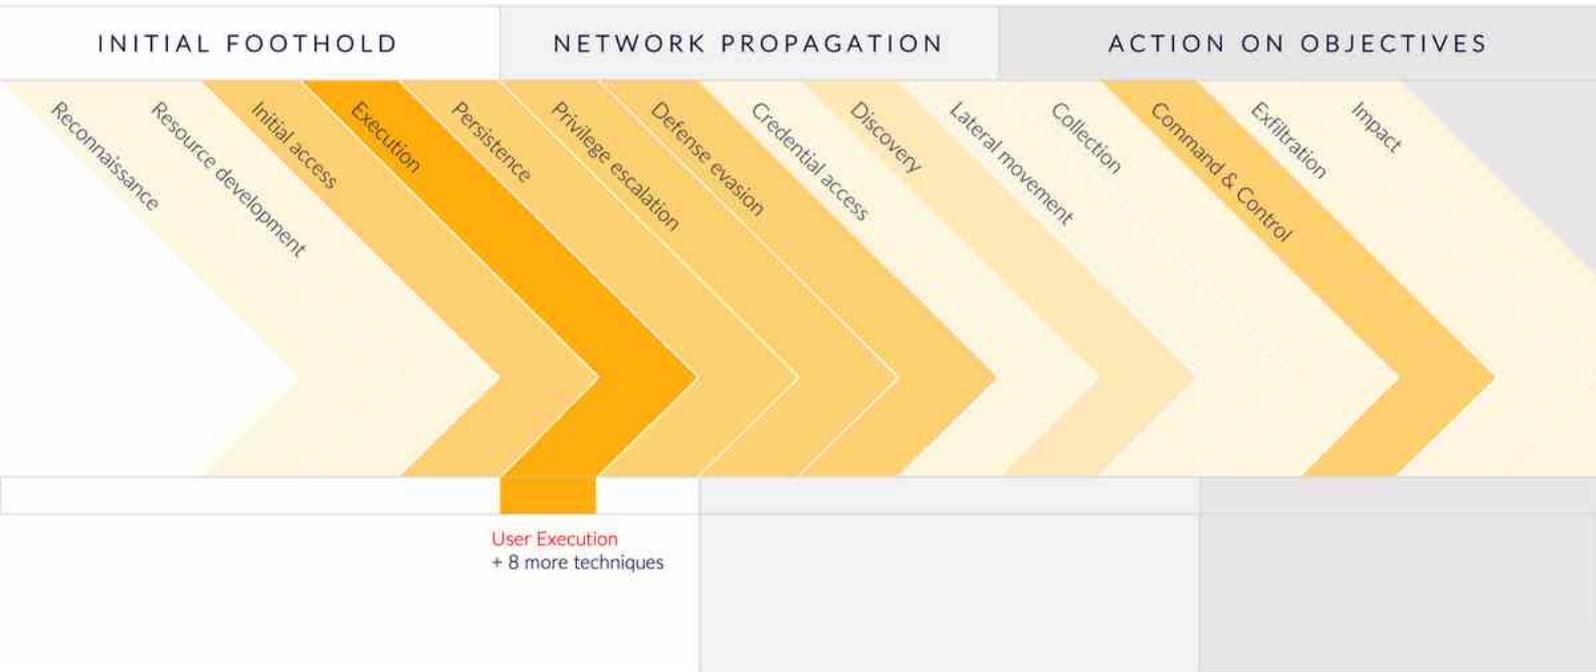
- Initial Foothold (49%)
- Network Propagation (33%)
- Action on Objectives (18%)

A third (~33%) of detections clustered in the post-compromise Network Propagation phase, constituting Defense Evasion (11%), Privilege Escalation (10.5%), Discovery (5.5%), Credential Access (3.86%), and Lateral Movement (1.87%).

This pattern underlines the importance of end-to-end visibility. Once attackers get in, much of the battle is fought in the middle of the chain, where they escalate access, evade controls, and move laterally long before triggering an obvious “impact” event.

MITRE ATT&CK heatmap

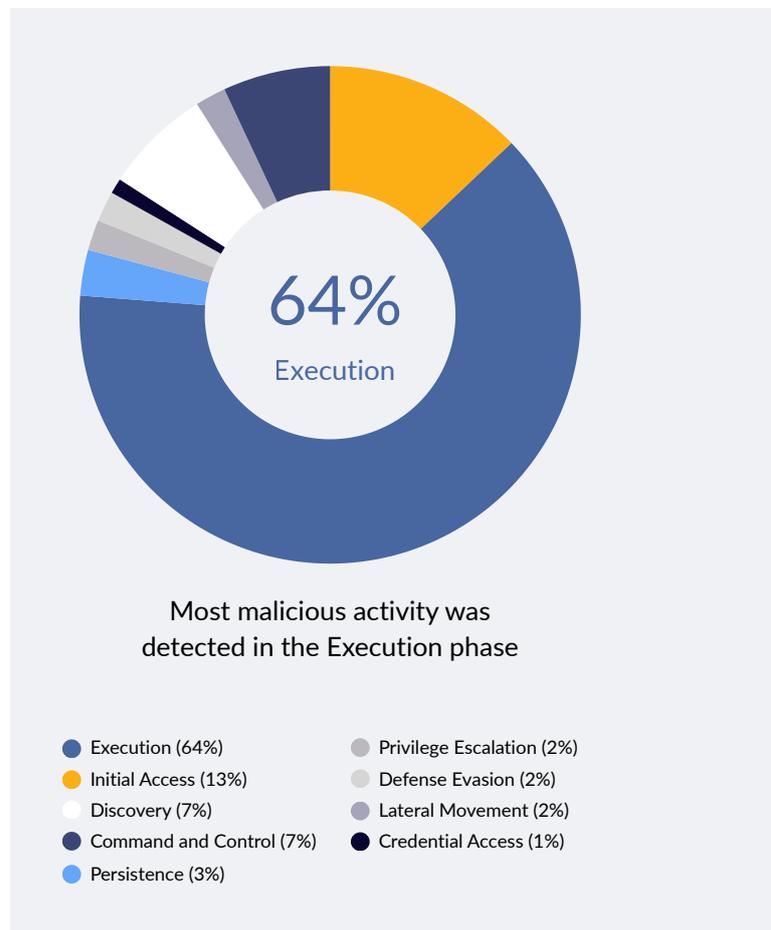
BASED ON 2025 SOC DATA



Distribution of malicious alerts

Of the detections mapped here, just 0.6% were deemed malicious after further investigation. The majority of these events were caught during the Execution phase (64%), followed by Initial Access (13%), Discovery (7%), and Command and Control (7%). The concentration of detections in the Execution and Initial Access phases indicates that, in most cases, malicious activity was identified before attackers could establish a deeper foothold.

However, the spread of detections across later phases such as Discovery and Command and Control reflects something equally important: full visibility across the entire attack path. Consistent with our assume breach approach, our detection strategy is designed to track attacker activity at every stage of the attack chain, not just at the point of entry. This gives us the best chance of stopping threats early, and in some cases allows us to observe attacker behavior in a controlled manner to understand their methods and access before expelling them from the environment.



Ten most common techniques

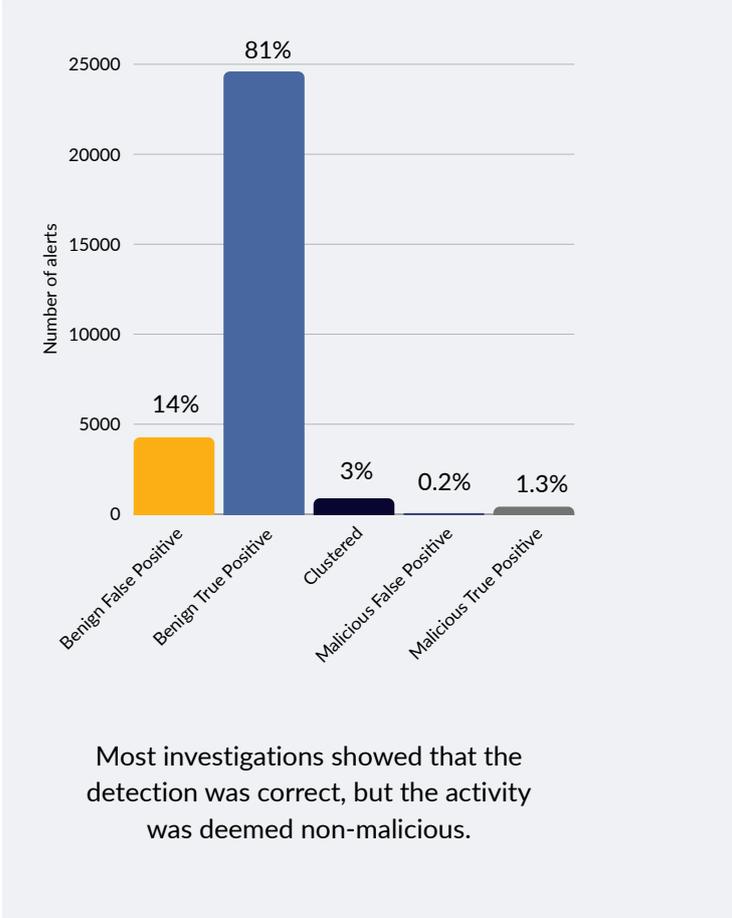
BASED ON 2025 SOC DATA

Technique	Tactic	Description	MITRE ID
User Execution	Execution	Adversaries trick users into executing malicious content, often through phishing or infected files.	T1204
Exploit Public-Facing Application	Initial Access	Adversaries exploit weaknesses in internet-facing applications or services to gain initial access.	T1190
Remote Access Tools	Command and Control	Attackers use legitimate remote access tools to control victim systems covertly.	T1219
Ingress Tool Transfer	Command and Control	Attackers transfer malicious tools or files into a compromised network to support further operations.	T1105
Impair Defenses	Defense Evasion	Adversaries disable or modify security software and controls to hinder detection and analysis efforts.	T1562
Command and Scripting Interpreter	Execution	Adversaries abuse command and script interpreters to execute commands, scripts, or binaries on a system.	T1059
System Binary Proxy Execution	Defense Evasion	Attackers misuse trusted, signed binaries to run malicious payloads, evading process and/or signature-based defenses.	T1218
OS Credential Dumping	Credential Access	Adversaries dump credentials to obtain account login and credential material, normally in the form of a hash or a clear text password, enabling lateral movement.	T1003
Event Triggered Execution	Privilege Escalation, Persistence	Attackers establish persistence and/or elevate privileges by configuring malicious code to run in response to specific system events.	T1546
Boot or Logon Autostart Execution	Privilege Escalation, Persistence	Adversaries configure system settings to automatically execute a program during system boot or logon to maintain persistence or gain higher-level privileges on compromised systems.	T1547

Signal-to-noise ratio

Our SOC data reflects some of the fundamental challenges facing the broader cybersecurity industry, namely managing alert noise and balancing speed with investigative depth. Every SOC must contend with significant noise. Detection rules need to be broad enough to catch subtle signs of intrusion alongside obvious ones, but this means benign activity inevitably gets flagged too. In 2025, 81% of security investigations were classified as benign true positives - correct detections that posed no real threat. Just 1.5% represented genuine malicious activity.

These numbers reflect a deliberate choice about sensitivity settings. An organization can easily reduce noise by raising its detection thresholds, but in doing so, it directly increases the chance that malicious activity will be missed. By maintaining relatively low thresholds, we accept more noise in exchange for catching a wider range of threats. This ratio is continuously refined by tuning and optimizing our rulesets to reduce false positives without compromising signal.



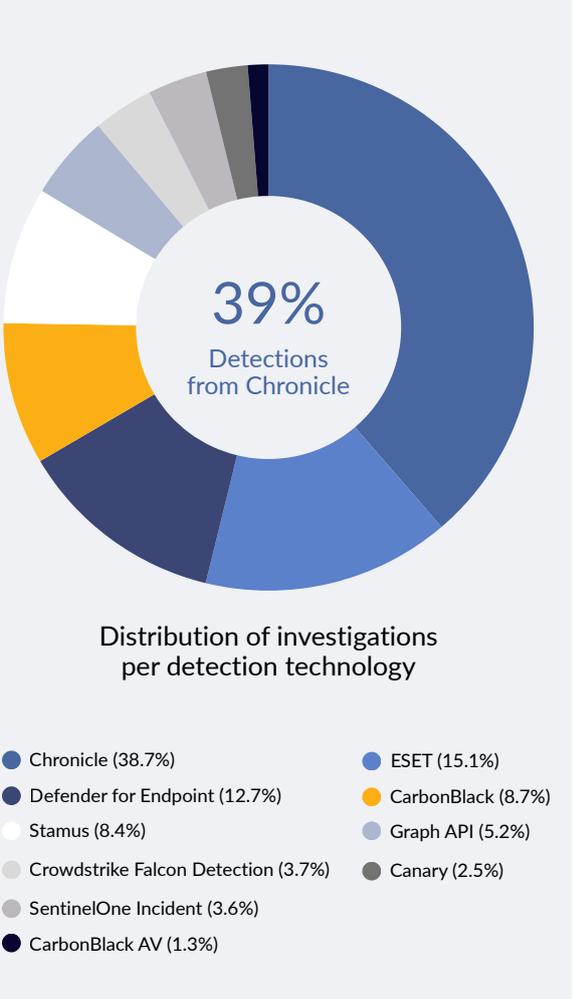
Glossary of terms

Benign True Positive	An alert that accurately matches detection logic but, upon investigation, is identified as legitimate, non-malicious activity.
Benign False Positive	An alert that does not match detection logic and is identified as legitimate, non-malicious activity.
Malicious True Positive	An alert that accurately matches detection logic and is identified as malicious activity.
Malicious False Positive	An alert that does not match detection logic but is identified as malicious activity.
Clustered	Related alerts that are grouped together for analysis purposes.

Diversity of detection technologies

To ensure sufficient visibility across various attack surfaces and kill chain phases, our SOC leverages a wide range of detection technologies, including Endpoint Detection and Response (EDR), Network Detection and Response (NDR), and deception technologies such as canaries. This is based on the understanding that no tool will catch everything and is central to our *detection in depth* approach.

In 2025, the majority of detections were triggered by Chronicle (38.7%), which represents our own proprietary ruleset.¹ Of the third-party technologies we employ, the highest share of detections came from ESET (EPP/ EDR) at 15.1%, followed by Defender for Endpoint (EDR/XDR) at 12.7%.



^[1] Chronicle detections (38.7%) reflect alerts triggered by our proprietary ruleset only; this figure does not represent the full scope of custom detections employed across our environment. Custom detection rules are also applied to Stamus and Carbon Black.

Speed matters, but depth determines outcomes

The SOC's core challenge isn't just to identify malicious events among the noise, but to do so quickly. Service Level Agreements (SLAs) define upfront how fast certain investigations must be completed, acting as an important benchmark. In 2025, our SOC closed 99.4% of investigations within the stipulated timeframe. Of that small fraction (0.6%) that exceeded it, nearly half were resolved within five (additional) minutes.

99.4 %

SOC investigations closed within the SLA window

Our philosophy is that SLA adherence is a measure of efficiency, not a proxy for quality. Some investigations require deeper analysis than the SLA allows, creating an inherent tension between speed and thoroughness. Getting this balance right often determines whether you're just detecting threats or actually responding to them effectively.

We therefore consider it a necessary trade-off to exceed the SLA for complex threats requiring in-depth investigation, as our SOC's function extends beyond initial triage. We use these metrics not as a quality score, but as a tool to identify opportunities for automation, with the overall goal of allocating analyst time and expertise to the most complex and high-impact threats.

From detection to investigation

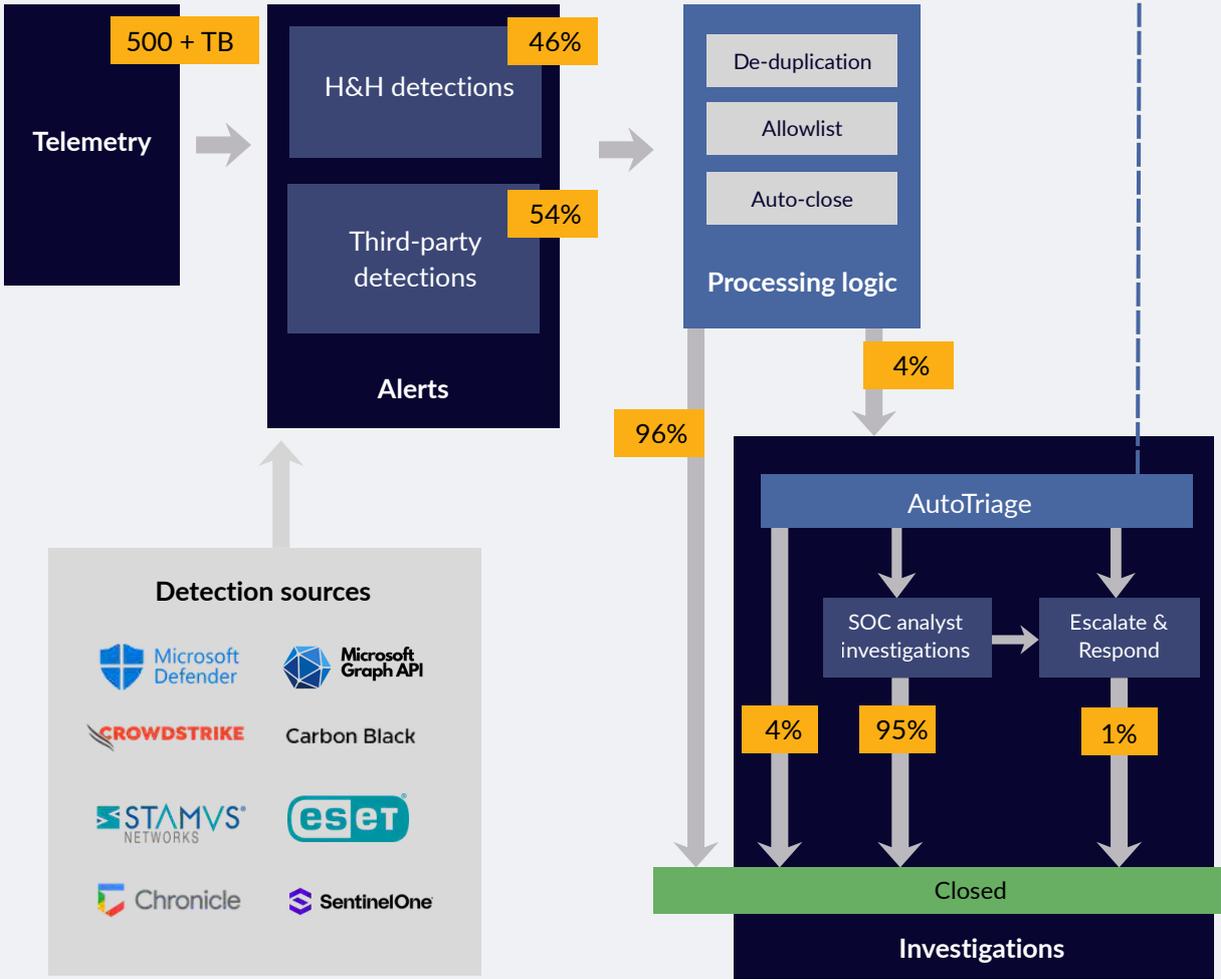
As seen in the image below, the alerts handled by the SOC represent just a small slice of the total alerts generated before processing logic is applied (~4%). This total alert volume is generated by Hunt & Hackett's own proprietary ruleset, as well as third-party detection technologies. Alerts are first assessed by processing logic to determine whether they meet the criteria for investigation, then filtered through our self-developed AutoTriage framework. What remains is passed to the SOC, forming the basis of our security investigations.

Typical alert flow

Based on a subset of 2025 SOC data

AutoTriage is our proprietary framework for automatically handling routine, non-malicious security alerts. It evaluates each alert against predefined conditions and closes it if benign.

[Learn more.](#)



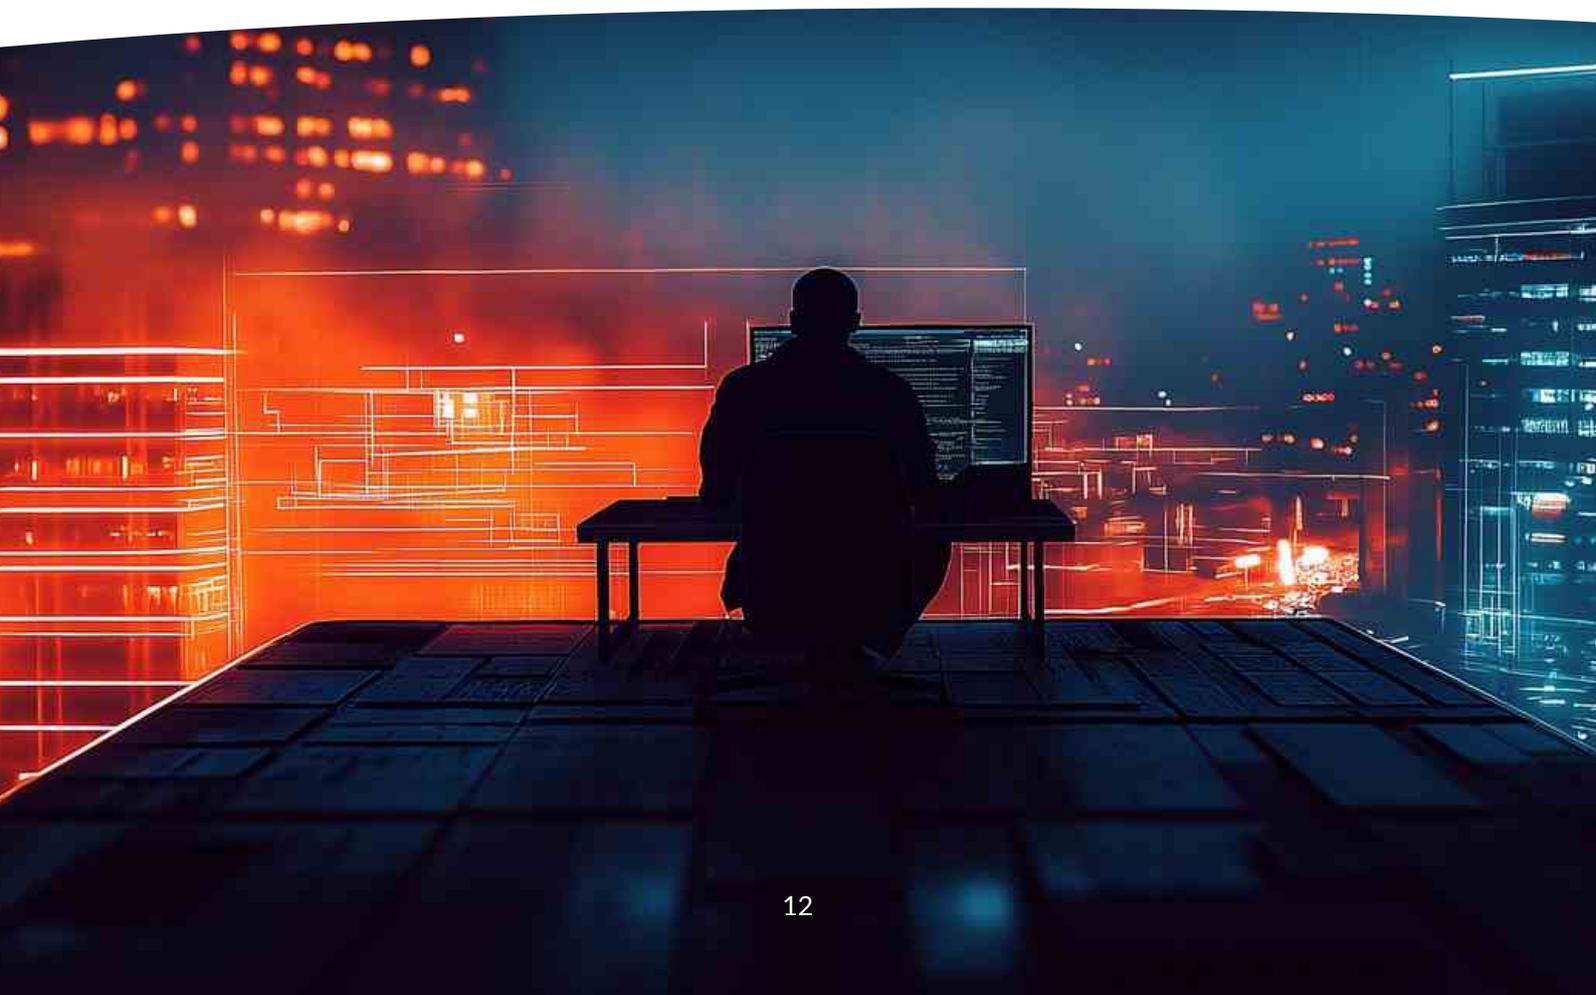
Platform availability as a prerequisite

Finally, the availability of our security platform - encompassing the SIEM, SOAR, and data pipelines responsible for ingesting and analysing security data - underpins every detection and response action the SOC performs. We guarantee 99% platform availability for all customers. While this is a conservative contractual commitment, in practice we achieve significantly higher uptime, typically 99.99% or better.

This level of platform reliability is an important, but often overlooked prerequisite for continuous detection, as even brief periods of downtime can create critical visibility gaps for an attacker to exploit.

99.99 %

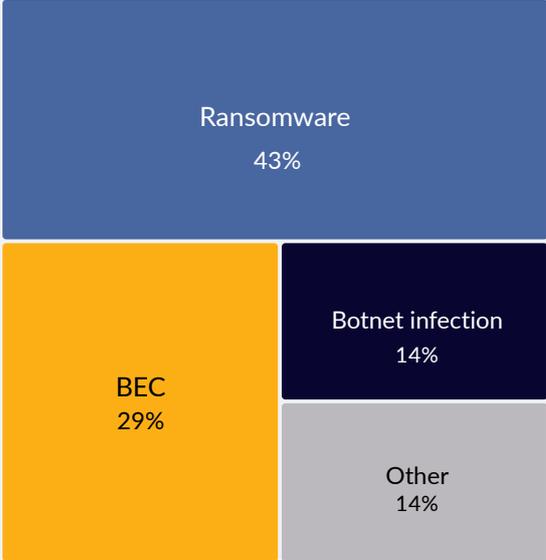
Availability of the security platform



Insights from Incident Response engagements

Our incident handlers responded to a variety of cases in 2025, with financial gain being the dominant motive in 71% of engagements. Within this category, ransomware was the most common incident type we handled at 43%, followed by Business Email Compromise (BEC) at 29%.

Initial access vectors varied across incidents. Remote services were a common entry point, often through vulnerable management software or edge devices. In other cases, attackers used valid credentials (typically obtained through phishing or password reuse) to access target networks. There was a strong emphasis on credential and Active Directory exploitation, including successful Kerberoasting attacks to extract service account credentials, exploitation of weak or easily crackable passwords, and AiTM (Adversary-in-the-Middle) phishing campaigns to bypass MFA for cloud-based accounts.



Distribution of incident types handled by Hunt & Hackett's IR team in 2025.

Note: Details of cases labeled 'other' are not disclosed to protect client confidentiality.

Comparing our view with wider trends

The prevalence of ransomware in our caseload is worth noting, as it runs counter to a broader global trend where BEC is now often cited as the more frequent attack vector.^[1] We believe this divergence stems from the specific risk profile and market segment of our client base, as well as the relatively small sample size we are drawing from.

It's good to note that the type of incidents an organization sees often depends on its specific role in the market. As an MDR and IR provider, our data is weighted towards technically complex intrusions that require specialist intervention, such as ransomware. Other stakeholders with a different market position, such as legal or insurance parties, may see a different cross-section of cybercrime. This is highlighted by Erik Jonkman, incident responder and breach counsel at Kennedy van der Laan, who noted the dominance of BEC in his 2025 caseload: "In terms of the cases we've been seeing most often, Business Email Compromise has been leading the charge, followed by ransomware and data breaches caused by human error."

For attackers, BEC is attractive because it is less technically complex than ransomware, lower risk, and has a higher success rate. Individual BEC cases typically result in smaller financial losses than the average ransom demand, but the simpler technical requirements mean attackers can run these campaigns at scale. The result is that BEC is now the leading form of cybercrime in terms of total global financial losses.^[2]

Known and detectable

Across these investigations, we noticed a recurring pattern: attacker techniques, though sometimes sophisticated, were widely known and detectable in principle. EDR and other tooling could have caught most of them, but critical gaps prevented this. Telemetry was often incomplete, with missing audit logs, network data, or insufficient log retention. In some cases, targeted systems fell outside the security scope entirely, providing attackers with convenient blind spots. These factors hindered detection and investigation in 86% of cases.

In some cases, a lack of fundamental security hygiene enabled or compounded attacks. Structural issues such as outdated systems, unpatched vulnerabilities and over-permissive access controls either opened the door to attackers or allowed them to move deeper into targeted networks. According to Erik Jonkman, this remains a common issue.

"Within the cyber community, it is obviously more interesting to focus on the advanced tactics used by leading threat actors. Nevertheless, I should emphasize that over 95% of the incidents we work on are simply caused by a lack of basic cyber hygiene. And while many organizations have invested in becoming more cyber resilient, the inconvenient truth is that society as a whole is still in the jungle, with enough low-hanging fruit for cybercriminals to live on for years."

Erik Jonkman
Partner at Kennedy Van der Laan

At Hunt & Hackett, we see the issue of poor security hygiene not as a simple failure of will, but as a symptom of a deeper execution challenge. As our co-founder Ronald Prins suggests, many organizations know they should fix the "low-hanging fruit" but struggle to do so consistently across complex environments shaped by technical debt and competing priorities.

"We've been talking about fixing low hanging fruit for years. The reality is that a large share of today's attacks remain preventable through cyber hygiene and limiting technical debt. If we're honest, the problem may not be awareness, but capability. And if that's the case, our focus should shift from chasing perfect prevention to detection environments that function reliably in the real world."

Ronald Prins
Co-founder at Hunt & Hackett

This capability gap is often compounded by a 'tools-based' security mentality, which centers around the assumption that deploying a single solution, like Endpoint Detection and Response (EDR), will provide sufficient and uniform coverage.

"What often happens is that organizations install an EDR and think it will catch everything, but our data shows that this is not the case. To protect an organization, you need visibility over the entire attack path and detection rules that reflect the current threat landscape."

Ronald Prins
Co-founder at Hunt & Hackett

Gaining visibility across the full attack path requires integrating a broad set of telemetry sources, including OS logs, EDR telemetry, network and identity signals, firewall data, and more. In many MDR setups, telemetry is treated primarily as input for detection and response. More mature programs treat it as a strategic asset: using attack graphs, control mapping, and continuous validation, telemetry can be translated into actionable insight that prioritizes hardening and closes control gaps - improving prevention, detection, and response while reducing technical debt over time.

When a detection or incident exposes a weakness, it presents an opportunity to refine security controls so the same gap can't be exploited again. Over time, this iterative cycle of detection, analysis, and remediation forms the basis of an adaptive security posture.

Beyond this, organizations should develop a thorough understanding of the attack paths that exist within their own environments. Trust relationships, access permissions, and system interdependencies all create routes that an attacker could follow to reach critical assets. By mapping these paths, security teams can make more informed decisions about where to concentrate hardening and monitoring efforts.

Securing a modern organization is inherently complex, and attempting to address every gap simultaneously is neither realistic nor effective. A more pragmatic approach starts with understanding the environment, identifying the most critical exposure points and addressing those first. Once a solid baseline has been established, organizations are in a better position to tackle the more complex weaknesses tied to technical debt and accumulated system dependencies.

EDR variability in practice

DEPLOYMENT ≠ UNIFORM OUTCOME

Over the past year, our Breach & Attack Simulation (BAS) testing repeatedly showed that Endpoint Detection & Response (EDR) can behave inconsistently within the same environment. Endpoints running identical agent versions and policies did not always produce identical detection, telemetry, or blocking outcomes - sometimes seemingly influenced by where a device sits in the network.

In several cases, expected detections or blocks occurred on most systems but failed on a subset, without any obvious difference in agent versioning or configuration. This suggests real-world EDR outcomes can be affected by environmental factors such as policy drift, sensor health, resource constraints, or differences in the telemetry path. In other words, deployment evidence is not the same as performance evidence. Regular validation across segments and endpoint archetypes is essential to confirm controls work as intended in your specific environment.

This challenges compliance-style assumptions that EDR deployment equals uniform protection. The practical takeaway is that control performance must be validated, not assumed. Continuous, environment-specific testing helps uncover “quiet” pockets where visibility or prevention is weaker, before an attacker finds them.

Looking ahead, this issue may become more pronounced as EDRs increasingly incorporate AI-driven detection and response. Greater model complexity can amplify differences in outcomes across environments, making ongoing validation even more important.





CHAPTER TWO

The evolving attacker playbook

In 2025, threat actors refined rather than reinvented their playbooks, combining familiar tactics with newer methods that are proving effective. The fundamentals may not have changed all that much, but it is clear that attackers have become more systematic at finding and exploiting organizational weaknesses. Technical debt, growing cloud adoption, and complex supply-chain dependencies continue to widen the attack surface, while unpatched edge devices and fragile identity controls serve as reliable entry points. This chapter examines how these developments are shaping the attacker playbook and what they mean for defensive strategies in 2026.

Key trends

- **Cybercriminals as technical debt collectors:** Attackers are systematically exploiting security gaps created by technical debt across legacy systems, dependency chains, and layered infrastructures, treating these accumulated weaknesses as a monetizable opportunity.
- **The expanding attack surface:** The number of exploitable entry points continues to grow as threat actors aggressively target vulnerable network edge devices, abuse valid accounts in cloud environments, exploit complex third-party supply chain dependencies, and take advantage of poorly governed AI integrations.
- **Identity as a leading attack vector:** Attackers are increasingly bypassing traditional defenses by logging in rather than breaking in, leveraging stolen credentials and session tokens to obtain and maintain access through legitimate identities.
- **AI widens attack-defense asymmetry:** Generative AI is accelerating the scale and efficiency of cyberattacks, exacerbating the already asymmetric relationship between attackers and defenders. While novel methods are beginning to emerge, the full impact on the threat landscape remains to be seen.

Paying the interest on technical debt

Today's attackers are increasingly operating as technical debt collectors, systematically exploiting the security gaps created by years of 'we'll fix it later' decisions made during development.^[3]

This 'debt' refers to the accumulated cost of prioritizing speed and immediate functionality over long-term quality and maintainability. While often necessary for maintaining forward momentum, these pragmatic trade-offs quietly compound across an organization's infrastructure, steadily expanding the attack surface over time.^[4] In 2026, attackers are viewing this accumulated debt as an opportunity.

"Ransomware groups are increasingly operating as technical debt collectors. Once it's there, there's a good chance it's being exploited. Attackers view their own efforts as a (non-governmental) tax on technical debt."

Ronald Prins
Co-founder at Hunt & Hackett

Despite how widespread this issue is (most companies have at least some amount of technical debt), its security impact remains under-recognized. Part of the problem is that technical debt is not always easy to see and much harder to fix. This debt accumulates at scale through legacy systems, embedded components, and dependency chains, much of which is poorly inventoried or only partially understood. The layered nature of the problem creates complexity, particularly in large organizations with complex infrastructures.

As Jurjen Harskamp notes: "You want to be sure all the windows and doors are closed in your environment, but in large organizations with complex IT and OT estates, that's extremely difficult. Technical debt isn't just at the application level, but the underlying libraries, services, and vendor components it depends on. Because systems are layered, updating a single vulnerable dependency can have a cascading impact across multiple applications (as the Log4Shell episode highlighted), so remediation is rarely a simple patch-and-move-on."

"That's what attackers exploit - the gaps and delays created by technical debt that can't be resolved quickly without operational risk."

Jurjen Harskamp
CEO at Hunt & Hackett

For financially motivated cybercriminals, the challenge isn't just to identify exploitable technical debt, but to monetize it. Importantly, extortion in these cases no longer always looks like ransomware. Increasingly, threat actors are shifting away from encryption-based tactics and instead opting for "data-only" extortion, where attackers exfiltrate sensitive information and threaten publication or customer notification.^[5] In regulated sectors, this dynamic can be even more coercive, as the attacker's leverage is not only reputational but also tied to potential legal exposure and mandatory reporting obligations.

Ultimately, technical debt is not an isolated IT problem, but a foundational security issue with no easy answers. While addressing it is a complex undertaking, it is also a clear imperative for building long-term resilience. In the following section, we take a closer look at one of the direct consequences of this debt - a wider and more fragile attack surface.

The expanding attack surface

The modern attack surface is in a state of constant expansion, driven by accelerating cloud adoption, complex supply chains, IT/OT convergence, and the rapid (and often ungoverned) integration of new technologies like AI.

While this expansion is a natural byproduct of growth and innovation, the presence of underlying technical debt can turn these surfaces into exploitable weaknesses. The result is a landscape of new and poorly understood entry points that attackers are learning to exploit with increasing efficiency.

Edge device exploitation

In recent years, threat actors have increasingly exploited network edge devices such as routers, IoT devices, firewalls, cameras and sensors, to gain initial access into corporate environments. These devices are attractive targets because they sit at the gateway to an organization's internal network and are often unpatched, poorly monitored, and directly exposed to the internet. Critical vulnerabilities in these devices have become an almost weekly occurrence, with companies like Cisco, Palo Alto Networks, Ivanti and Fortinet having made repeated disclosures over the past year.

In 2025, Verizon found that 22% of all exploited vulnerabilities targeted VPN and edge devices, almost eight times higher than the previous year. The median time to mass-exploit a vulnerability on an edge device was zero days from its CVE publication. And of 17 sampled edge device CVEs, nine were added to the CISA KEV catalog on or before their CVE publication date. This means that attackers were exploiting these vulnerabilities before they were even publicly documented.^[6]

By contrast, the median time to patch known vulnerabilities was 32 days, with many devices going unpatched for much longer. According to Verizon, only 54% of vulnerable edge devices were fully remediated last year, effectively leaving the door open for attackers.^[6]

“Exploitation of edge devices, particularly vulnerable VPNs and firewalls, has been a common theme in 2025. The majority of the ransomware incidents we have been involved in concerned unpatched vulnerabilities caused by inadequate patch management.”

Erik Jonkman

Partner at Kennedy Van der Laan

Exploitation of edge devices has become a staple in the playbooks of both cybercriminals and state-sponsored APTs. A recent Amazon report found that Russian GRU-linked actors have pivoted from traditional vulnerability exploitation to primarily targeting edge devices in campaigns against Western critical national infrastructure.^[7] Similarly, Chinese groups like Salt Typhoon and RedNovember have actively exploited vulnerabilities in devices from Cisco, Ivanti, Palo Alto Networks and Juniper for global espionage, with Salt Typhoon alone breaching 600 organizations across various critical sectors.^[8]

Edge devices are particularly vulnerable to exploitation because they often have weak default security settings, run outdated software or firmware, or use legacy hardware. They're also difficult to monitor, largely due to the high cost of device-specific configurations, variations in log details and forwarding capabilities, and the economics of data ingestion for MDR providers. By default, these devices are often excluded from regular Managed Detection and Response (MDR) services, with only a few providers accepting data from them.

“Edge devices often fall outside the standard MDR scope, which creates a meaningful detection blind spot. These data sources are typically high-volume and low-signal, and extracting value from them requires customized parsing, filtering, and detection logic - hence why they are frequently excluded from standardized MDR offerings. All this means that if threat actors compromise such systems they can be there unnoticed for a long time.”

Jurjen Harskamp
CEO, Hunt & Hackett

The lack of visibility this creates poses significant risks, particularly for organizations with operational technology (OT). This is because edge devices can serve as a bridge from IT into cyber-physical systems. OT environments are often difficult to patch, built on long lifecycle systems, and are operationally sensitive, meaning compensating controls and visibility matter as much as “fixing the vuln.”

Attackers who pre-position in edge devices can wait for the right moment to move laterally, map industrial networks, and target control systems indirectly. This is one reason why pre-positioning campaigns by advanced actors such as Salt and Volt Typhoon matter beyond espionage - in the wrong context, access can translate into disruption of physical services. For critical sectors, resilient monitoring must extend to the boundary devices and the pathways that connect IT to OT.

When visibility is limited, validate

Given the difficulties in detecting edge device exploitation, validation mechanisms like threat hunting act as an important safeguard. In 2025, we initiated a project to collect and enrich Indicators of Compromise (IOCs) relevant to our customers' threat landscapes, then used them to run automated threat hunts across customer environments for signs of compromise that may have slipped through.

This approach is grounded in the *assume breach* principle, which states that no defense is perfect, and some sophisticated threats will inevitably bypass prevention and detection controls. Retroactive threat hunting also helps us identify why past threats were missed, feeding improvements back into our defenses and creating a continuous feedback loop.

CASE STUDY

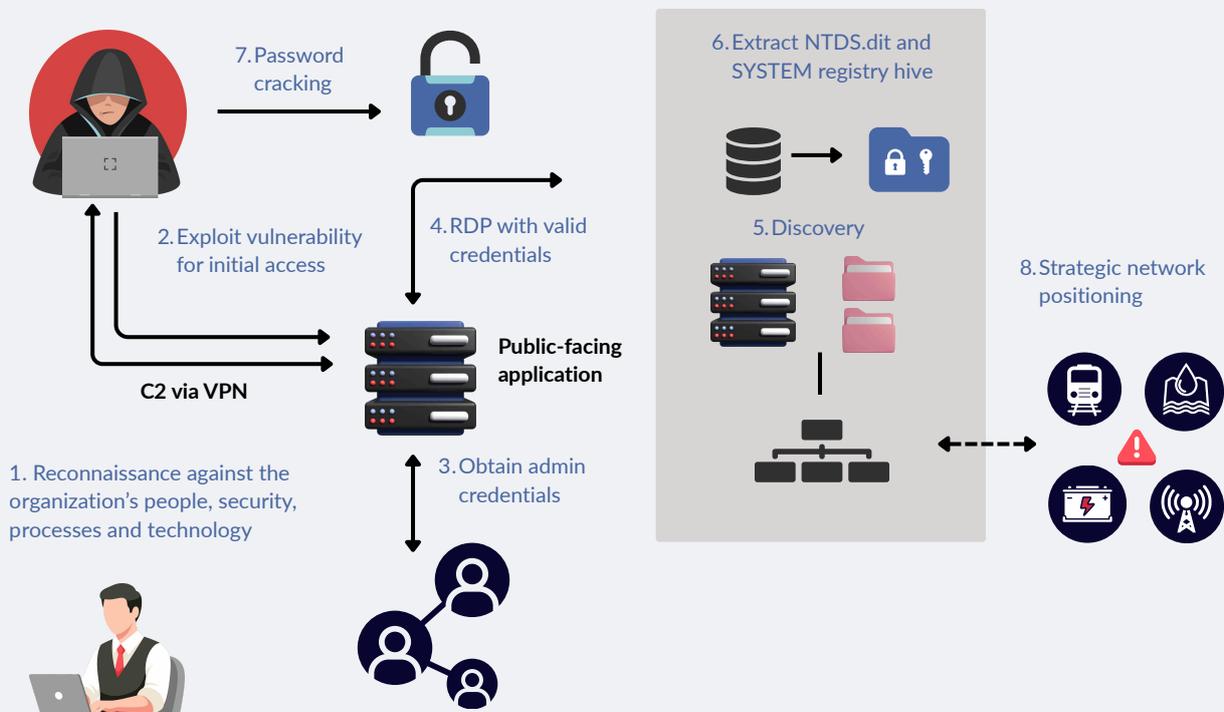
Edge devices open the door to OT networks

Few groups have operationalized this approach more effectively than Chinese APT Volt Typhoon. The group is known for systematically exploiting internet-facing edge devices from vendors like Fortinet, Ivanti, and Cisco, often using publicly available code for known vulnerabilities alongside zero-day exploits. This approach is a core part of its campaigns focused on pre-positioning within critical national infrastructure.

Once inside, Volt Typhoon shifts to a living-off-the-land approach, using valid credentials and native tools to blend in with normal network traffic and maintain long-term, hard-to-detect persistence. The group methodically builds full domain control and OT-adjacent access by repeatedly extracting the Active Directory database (NTDS.dit), cracking passwords offline, and using the stolen credentials to test access to domain-joined Operational Technology (OT) assets. This positions them to disrupt physical operations should geopolitical tensions escalate.^[9]

Typical Volt Typhoon attack flow

Source: [CISA](#)



RECONNAISSANCE

INITIAL ACCESS

LATERAL MOVEMENT

IMPACT

Cloud account abuse

As organizations continue migrating workloads to the cloud, attackers are following close behind. And rather than exploiting technical vulnerabilities, they're increasingly walking in through the front door, mostly unseen.

According to CrowdStrike, valid account abuse (T1078.004) has become the primary initial access vector for cloud intrusions, accounting for 35% of incidents in the first half of 2024.^[10] This trend ties directly to the surge in identity-centric tactics discussed later in this chapter. Valid account abuse enables multiple attack objectives at once, including initial access, lateral movement, persistence, privilege escalation, and defense evasion.^[11] It has been a feature of several high-profile attacks targeting cloud environments in the past year, most notably a wave of data theft incidents impacting Salesforce customers. The case study on the next page provides an overview of these cases.

“Data-rich CRM platforms like Salesforce have become high-value targets in financially motivated campaigns. By using social engineering to compromise legitimate accounts, attackers can operate under the guise of normal user activity - making detection significantly harder and giving them the time to exfiltrate large volumes of data for extortion.”

Jurjen Harskamp
CEO, Hunt & Hackett

The Salesforce incidents highlight two important aspects of modern cloud attacks, namely the exploitation of human trust (allowing attackers to access valid accounts and escalate privileges) and the manipulation of system trust (between integrated cloud apps). This reality demands a dual focus for defenders - hardening the human layer against social engineering and securing valid accounts, while also managing the inherent risks posed by a deeply interconnected digital ecosystem. The challenge is not just for an organization to secure its own accounts, but also to manage trusted relationships with its suppliers. As we show in the next section, this necessity extends far beyond the cloud.

CASE STUDY

Salesforce data breaches

In a widespread campaign spanning from at least May 2025 into early 2026, threat actors affiliated with the hacking collective Scattered LAPSUS\$ Hunters targeted dozens of high-profile Salesforce customers, including Google, Allianz Life, LVMH brands and Coca-Cola Europacific Partners. The attackers used a consistent playbook focused on social engineering to bypass security controls and exfiltrate sensitive CRM data for extortion.

The breaches typically began with attackers harvesting credentials through reused or phished single sign-on (SSO) logins, then using voice phishing (vishing) to impersonate IT support in phone calls with employees of the target organizations. During these calls, threat actors guided employees to visit Salesforce's *Connected App* setup page and tricked them into entering a connection code. This authorized a malicious, rebranded version of the Salesforce Data Loader, a proprietary tool used to bulk import data into Salesforce environments.

The malicious app generated long-lived OAuth tokens for persistent access, allowing the attackers to bypass multi-factor authentication and quietly exfiltrate large volumes of CRM data without triggering security alerts. In some cases, they also pivoted into connected platforms like Microsoft 365 or Okta to expand the scope of the compromise. After stealing the data, the attackers extorted victims by publicly naming impacted companies and leaking data samples, demanding ransom payments to prevent further leaks. Salesforce said its own systems were not directly compromised, and the incidents didn't stem from a known vulnerability in its technology.^[12]

The digital supply chain

The exploitation of digital supply chains is a defining feature of the modern, expanded attack surface. The popularity of this tactic has surged in recent years; according to Verizon, third-party involvement in breaches has doubled in the past year alone,^[13] and a World Economic Forum survey now reports that 65% of large companies see it as their greatest cybersecurity challenge.^[14]

During a supply chain attack, threat actors exploit trusted relationships between organizations. SaaS providers and open-source software libraries are particularly attractive targets because a single compromise can provide access to hundreds of downstream customers. As Jurjen Harskamp explains, threat actors are increasingly targeting systems that build and distribute software to achieve legitimacy at scale.

“Adversaries increasingly aim for the systems that build, sign, and distribute software, such as CI/CD pipelines, artifact registries, package managers, and code-signing keys. The reason is simple: compromise a single build or signing process and you can scale access across many organizations, often with a high degree of legitimacy.”

Jurjen Harskamp
CEO, Hunt & Hackett

At its core, this approach weaponizes trust. As Ronald Prins notes, targeting systems that build and sign software effectively turns a defender's own security controls against them.

“Signed binaries, trusted update channels, and ‘approved’ dependencies can turn the defender’s controls into the attacker’s delivery mechanism. This is also why CI/CD security is no longer ‘just a developer concern’, but a frontline control point for enterprise and national resilience.”

Ronald Prins
Co-founder, Hunt & Hackett

Acting on this knowledge requires a shift from implicit trust to explicit verification - a responsibility shared by both software providers and their customers. For providers, this involves hardening CI/CD pipelines, rigorously protecting code-signing keys, and securing software dependencies. For their downstream customers, it means continuously validating software integrity and, crucially, applying the principle of least privilege to ensure that even trusted applications run with only the minimum permissions necessary.

Ungoverned AI integrations

As AI moves from standalone tools into the fabric of everyday business systems, it is expanding the attack surface on multiple fronts. What started as a handful of experimental pilots has become a dense mesh of chatbots, copilots, embedded models and third-party integrations, each with its own data flows, dependencies and failure modes.

One of the most immediate risks this presents is unwanted information exposure. Employees increasingly turn to generative AI tools to draft emails, summarize documents or write code. Verizon notes that a significant share of staff now routinely access Gen-AI platforms from corporate devices, with most of them doing so via non-corporate email accounts.

In practice, this means sensitive content (contracts, source code, internal strategy) is being pasted into systems outside the organisation's control, often without any logging, policy or review.^[13]

Erik Jonkman notes that recent warnings about unintentional data exposure in this context are "absolutely justified".

"AI is being adopted at scale, but many organizations are still lagging behind when it comes to AI governance and controls to prevent this unwanted data exposure. I expect to see an increase in these types of incidents in 2026."

Erik Jonkman
Partner at Kennedy Van der Laan

At the same time, the AI supply chain has become a target in its own right. Modern AI stacks depend on large open-source ecosystems, cloud services and specialized frameworks. Malicious packages in repositories like PyPI, poisoned training data, and traditional software bugs (such as remote code execution) in model-serving infrastructure all provide attackers with familiar ways in, wrapped in unfamiliar branding.^[15]

Finally, as AI is integrated into operating systems, it is becoming a default feature of endpoints rather than an optional extra. Gen-AI assistants are being embedded into mobile and desktop OSes, searching local files, emails and chats to provide "helpful" answers.

In BYOD environments where personal and corporate data co-exist and security controls are weaker, these components create new risks for information exposure outside of traditional monitoring.^[13] Long-term, ungoverned AI integration risks expanding the attack surface across all layers of computing infrastructure.

"The race to embed AI in all components throughout the OSI layers causes exponential gaps throughout all layers of the attack surface, if unmanaged. From a security point of view, that could pose a serious threat if it can't be managed and contained."

Jurjen Harskamp
CEO, Hunt & Hackett

This is a rare occurrence in security. Typically, new technologies introduce risk incrementally, allowing defenders time to adapt. With AI, the race to embed it everywhere means that the associated risks proliferate across all layers at once.

Organizations that treat AI as "just another app" may discover that they have not added one new layer to their attack surface but fundamentally altered its structure in ways we are only beginning to understand.

Identity as a leading attack vector

Identity-based attacks have emerged as a leading vector enabling cybercrime, accounting for 30% of global intrusions according to IBM data from 2025. This means that for every third attack, a hacker is logging into a victim environment rather than breaking in. This shift is partly driven by the growing efficacy of modern host-based security tools, particularly endpoint detection and response (EDR) solutions. As EDRs have gotten better at catching persistent malware on endpoints, the old attack chain – phishing a user to deliver a malware payload and establish a backdoor – has become less reliable, forcing attackers to adapt.^[16]

Identity-based attacks leverage legitimate credentials and session tokens to establish access, then use those valid accounts for lateral movement. They're much harder to detect because, from a technical standpoint, everything looks legitimate. The attacker isn't triggering alerts by dropping malware, they're just logging in like any other user. From their perspective, why bother smashing a window and climbing inside when you can walk in through the front door with a key?

Attack Type	MITRE ATT&CK Tactic	MITRE ID
Business Email Compromise	Initial Access, Credential Access, Collection, Impact	T1566, T1110, T1621, T1078, T1114, T1098, T1119, T1657
Kerberoasting	Credential Access	T1558.003
MFA Fatigue	Credential Access	T1621
Cloud Account Takeover	Defense Evasion, Persistence, Privilege Escalation, Initial Access	T1078.004, T1528
Adversary-in-the-Middle	Credential Access, Collection	T1557

MITRE techniques associated with common identity attacks

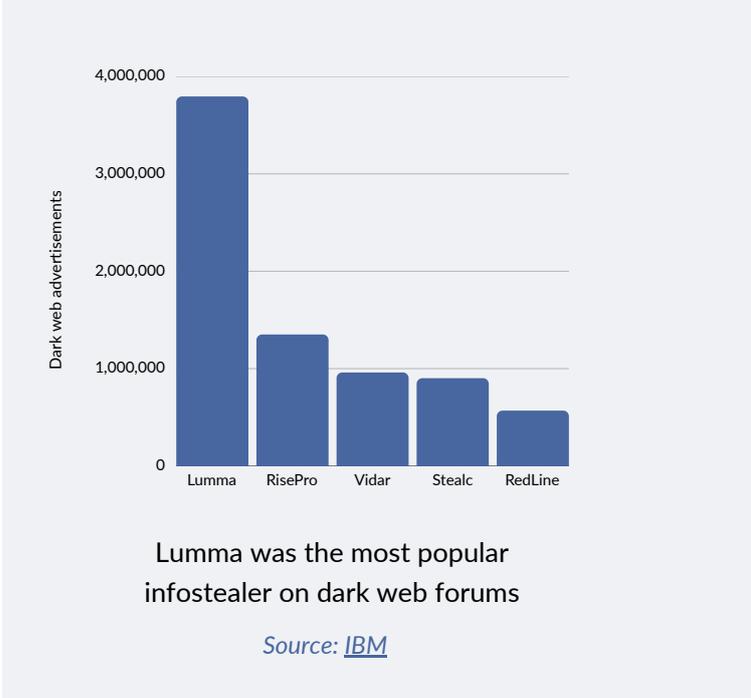
The rise of infostealers

Infostealer malware is one of the key enablers of identity-based attacks. As the name suggests, this type of malware is built to covertly extract sensitive data from victims. As identity has become a more popular attack vector, the number of infostealer infections has surged accordingly, rising 266% between 2023 and 2024.^[16]

Meanwhile, the number of infostealers delivered via phishing emails per week increased by 84% year-over-year.^[17] According to Hadrian, 64% of organizations globally had at least one known infostealer infection.^[18] Some of the most prominent strains seen in 2025 included Lumma, RisePro, Vidar, StealC, SnakeStealer and Redline.^[19]

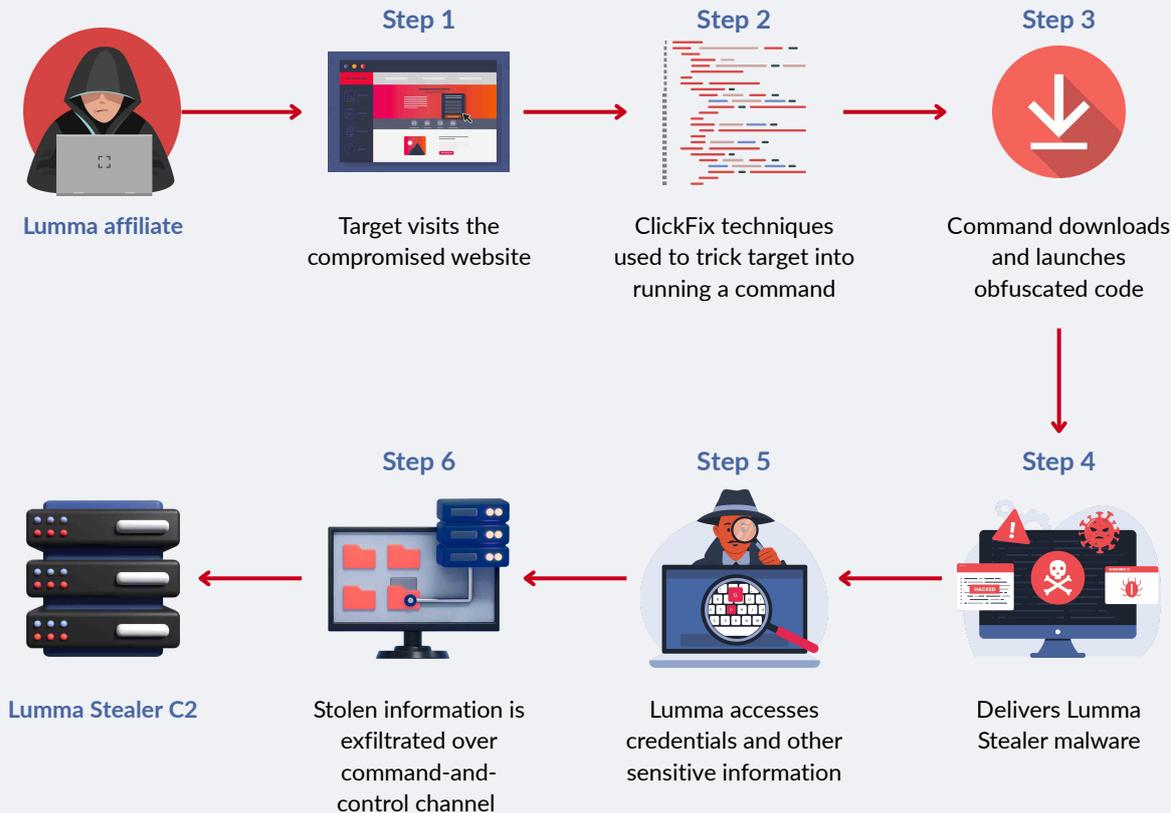
The real-world impact of this trend was demonstrated during a breach of Jaguar Land Rover (JLR) in March 2025, which preceded the large-scale incident that took place later in the year.

During this incident, HELLCAT ransomware group used infostealer malware to obtain the credentials of an LG Electronics employee, who had third-party access to JLR's Atlassian Jira server. The stolen credentials were used to access a Jira instance and exfiltrate large volumes of company data, resulting in more than 700 internal JLR documents being posted to the dark web.^[20]



Infostealer attack flow

Based on Lumma using ClickFix



CAPTCHA? Gotcha

As infostealers increased in popularity, attackers got more creative with their distribution methods. One notable example involved the delivery of infostealers through compromised WordPress websites with fake CAPTCHA and verification prompts embedded into the site. Victims were tricked into copying and pasting a malicious script and then running it under the guise of completing a legitimate verification step. This is a new social engineering tactic dubbed “ClickFix” by security researchers.

ClickFix was first observed by Proofpoint in March 2024 and quickly gained traction as a go-to technique for malware delivery.^[21] According to ESET data, the prevalence of this attack method grew 517% across 2025.^[22] One of the most impactful campaigns leveraging ClickFix was ClearFake, which distributed the Lumma and Vidar infostealers at scale, resulting in over 9,300 confirmed infections.^[23]

“We really saw ClickFix spread [last year]. It started being adopted by cybercrime actors and then nation state actors as well. It’s fairly common to see actors copying from each other like this.”

Martijn Grooten
Threat Intelligence analyst

Kerberoasting turns up the heat

In addition to the distribution of infostealers, Kerberoasting (T1558.003) has emerged as a popular technique in identity-based attacks. IBM observed a 100% increase in Kerberoasting during incident response engagements over the last year.^[24] For the uninitiated, Kerberoasting is a post-exploitation technique targeting the Kerberos authentication protocol, enabling adversaries to extract encrypted service account credentials from the Active Directory.

During such an attack, an adversary requests Kerberos service tickets for specific accounts in Active Directory, then attempts to crack (often poorly encrypted) tickets offline to obtain the underlying passwords. Once the plaintext credentials of the service account are obtained, the adversary can impersonate the account owner and inherit access to any systems, assets, or networks granted to the compromised account.^[25]

At Hunt & Hackett, we observed Kerberoasting in multiple incident response engagements throughout 2025. As with all identity-based attacks, Kerberoasting presents a challenge for defenders because it is not easily detectable. Most cybersecurity tools are not designed to monitor or analyze the behavior and activity of approved users, and because Kerberoasting does not rely on malware, traditional antivirus solutions are ineffective at detecting it. However, detection is possible when the right measures, such as custom detection rules, are in place.

Case Study

This was demonstrated by a recent incident handled by our Security Operations Centre (SOC). Our custom detection rules flagged possible credential abuse when a customer's service account suddenly generated an unusual, unprecedented volume of ticket requests. The source, an internet-exposed firewall, immediately drew our analysts' attention. At the same time, one of our honeypots triggered, indicating active exploration within the customer's network. Rapid analysis confirmed the device had been compromised and was operating as an internal host, enabling the SOC to contain the incident swiftly and block the attacker from moving deeper into the network.

Go vish

Finally, it's impossible to discuss identity abuse without acknowledging vishing, or "voice phishing" (T1566.004). This tactic directly supports identity-based intrusions, as attackers impersonate employees to manipulate corporate service desks into resetting passwords or enrolling new MFA devices, effectively bypassing verification controls altogether.^[26]

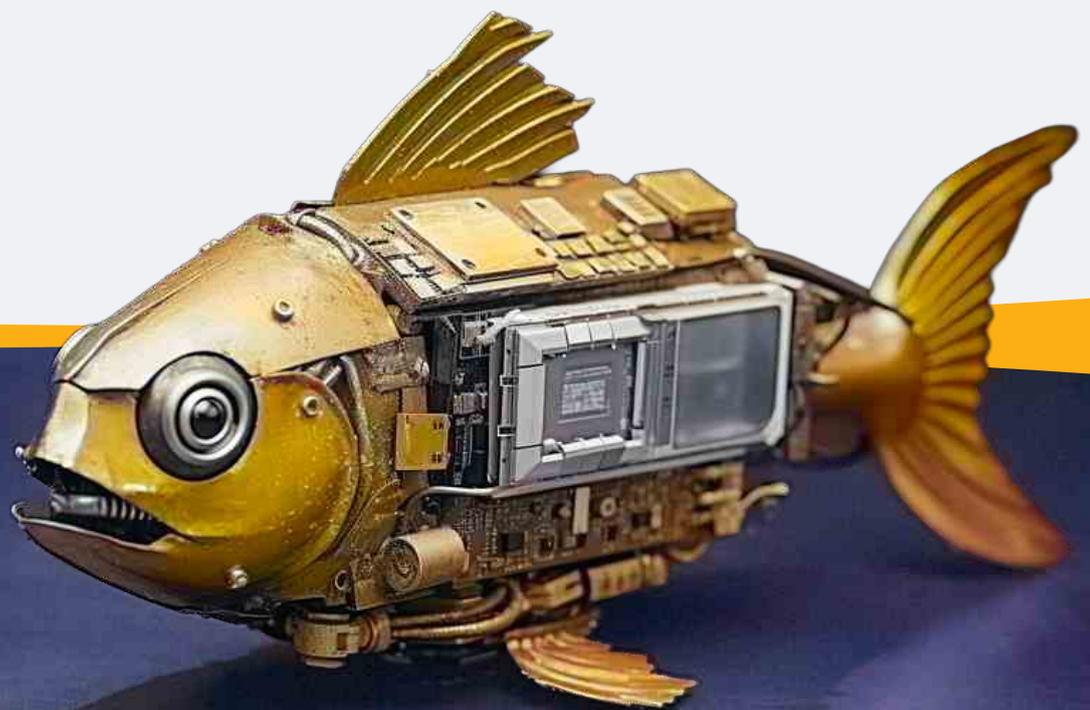
The widespread availability of AI tools has dramatically escalated the prevalence of vishing in recent years, with a reported 442% increase throughout 2024.^[27] This tactic has been broadly adopted by cybercrime actors and featured in several high-profile incidents, including Scattered Spider's attacks on UK retailers M&S, Co-Op, and Harrods, which reportedly caused hundreds of millions in financial losses.^[28]

Attackers are also taking steps to professionalize vishing beyond the use of AI. Over the past year, several cybercrime groups were observed openly recruiting callers on widely used underground forums. The advertisements typically favored native English speakers with knowledge of RMM tooling and experience conducting remote sessions.^[29]

"It's becoming increasingly normal for us to work remotely and conduct many of our interactions online. In that context, using deepfakes and vishing to commit crime is a very serious risk. Offenders are getting more professional, and their activities are becoming harder to recognize."

Daan Weggemans

Assistant Professor, Leiden University



Gen-AI exacerbates asymmetry between attack and defense

In last year's report, we observed that AI was helping attackers enhance the speed and scale of their operations, particularly in reconnaissance and social engineering. Although no truly transformative breakthroughs have emerged in the year since, it's clear that things are accelerating. The TTPs may be broadly the same, but attacks can now be orchestrated far more efficiently.

With that being said, the long-term potential of AI is still largely unclear. As Jurjen Harskamp, CEO of Hunt & Hackett, observes, "There are multiple realities at play. If you take a skeptical view, you're right, to a certain extent. If you take a very optimistic view, you're also right, to a certain extent. In that sense, everyone is currently right about AI, because there is still so much uncertainty. It's not unlike Schrödinger's cat: both alive and dead at the same time, until the box is opened and a definitive outcome is revealed."

Among the experts we consulted, there was broad agreement on one near-term effect: AI is contributing to an increase in both the frequency and intensity of cyber threats. By enabling attackers to operate at greater scale and efficiency, AI shifts the cost-benefit ratio and expands the pool of viable targets.

"In the past, when humans had to do most of the work, both attackers and defenders focused on big software systems and major vulnerabilities, because that's where the cost-benefit ratio made sense. With automation and AI, that changes. Many 'low-value' targets suddenly become attractive, because you don't need to start from scratch each time. You can take something that already exists and simply ask AI to adapt it to a new case, very quickly. So even a small company with 10 employees becomes a viable target - because for the attacker, it's all automated."

Stjepan Picek

Associate Professor, Radboud University

Jurjen Harskamp echoes this, noting: "The attack techniques might not have fundamentally changed, but the speed and scale at which they're executed is accelerating dramatically, which we're still quite unprepared for on the defensive side. The principle that you can only survive if you can operate at the speed of the attacker remains relevant and will become even more so."

AI is also starting to move the needle on the defensive side, especially in triage, enrichment, and correlation across high-volume telemetry. We've made a deliberate choice to avoid "autonomous SOC analyst" narratives in this report, but speaking generally, the limiting factors remain data quality, false positives, governance, and verifiability. During real incidents, decisions must be explainable, testable, and defensible.

How attackers used Gen-AI in 2025

Over the past year, attackers began deploying AI-enabled malware in live operations. Google has reported on five such malware strains: FRUITSHELL; PROMPTFLUX; PROMPTLOCK; PROMPTSTEAL; QUIETVAULT. Arguably the most interesting is PROMPTFLUX, an experimental dropper malware which interacts with Gemini's API to request obfuscation and evasion techniques mid-execution, enabling "just-in-time" self-modification to evade static detection. Malware with the ability to shapeshift mid-attack provides a significant new challenge for defenders.^[30]

Tactic	Technique	MITRE ID	Description
Execution	Command and Scripting Interpreter	T1059.005	Uses VBScript and Powershell for script-based execution.
Defense Evasion	Obfuscated/Encrypted File or Information	T1027	Dynamically rewrites code to evade detection.
Command and Control	Application Layer Protocol	T1071.01	Leverages HTTPS to communicate with Gemini AI endpoints.
Persistence	Scheduled Task/Job	T1053.005	Establishes persistence via scheduled tasks to maintain evolution cycles.
Discovery	System Information Discovery	T1082	Gather local system info for tailored payload adaptation.

MITRE ATT&CK mapping for PromptFlux | Source: [Medium](#).

Anthropic also weighed in on AI-assisted cyber operations. In November, the company claimed to have disrupted a "highly sophisticated" espionage campaign attributed to a Chinese state-sponsored group, where Claude Code was allegedly used to execute 80-90% of tactical operations autonomously. Anthropic described this as "the first documented case of a cyberattack largely executed without human intervention at scale," marking a significant advancement in weaponized AI.^[31]

However, the wider cybersecurity community quickly became critical of the report's framing. Although Anthropic spoke extensively about how AI can be leveraged to support intrusion activity, they failed to provide any IOCs, attribution details, or verifiable information. Many experts questioned how representative of reality the claims actually were.

Separately, CrowdStrike research uncovered a more subtle AI-related risk. Their testing showed that DeepSeek-R1, a Chinese LLM launched in January 2025, produces code with severe security vulnerabilities up to 50% more often when prompted with keywords or topics that are politically sensitive to the Chinese Community Party (CCP).

"The main underlying issue is that AI is not designed to be secure."

Stjepan Picsek
Associate Professor, Radboud University

In one example, mentioning “Tibet” caused the prevalence of vulnerable code to rise from 19% (baseline) to 27%. References to “Falun Gong” and “Uyghurs” produced similar results.

This may or may not be indicative of malicious intent - researchers have proposed an “emergent misalignment” theory, suggesting that pro-CCP training inadvertently taught the model to associate certain politically charged terms with negative characteristics, degrading output quality. Regardless of the cause, this introduces a new and subtle vulnerability surface for organizations relying on AI coding assistants.^[32]

A widening asymmetry

The relationship between attackers and defenders has always been inherently asymmetric. Defenders must guard against every possible attack vector, while attackers need only find one that works. AI doesn't change this fundamental imbalance. In many ways, it amplifies it.

While AI will likely lead to advantages on both the offensive and defensive sides, it appears that the short-term advantage currently lies with the attackers.

“Attackers have fewer obstacles to AI implementation than defenders. They can test new tools quickly, with little to lose. Defenders have to think about core functionality, responsible implementation, data security, and governance.”

Ronald Prins
Co-founder, Hunt & Hackett

Jurjen Harskamp echoes this: “The scale is structurally asymmetric in favor of attackers. While defenders must secure complex IT and OT environments continuously, attackers only need a single successful path for a limited time, using disposable infrastructure that can be quickly rebuilt. This asymmetry suggests that the first significant gains from AI in cybersecurity are more likely to benefit attackers - through scale, variation, and evasion - before comparable defensive advantages emerge.”

“I think at this point, AI is causing more problems than it is providing solutions in terms of cybersecurity.”

Daan Weggemans
Assistant Professor, Leiden University

The honest answer is that it is still too early to determine the lasting impact of AI on cybersecurity. Defenders are improving their ability to leverage AI to process large volumes of data more efficiently, which is an essential counterweight to the offensive side. However, accuracy and reliability remain significant challenges.

The technology is evolving faster than our collective ability to fully understand its implications, and that uncertainty cuts both ways. The only prediction we are confident in is that, by the time we write the 2027 edition of this report, many of today's assumptions (on both offense and defense) will already be outdated.



CHAPTER THREE

Cyber at the geopolitical frontlines

In last year's report, we observed that cyber operations were increasingly being integrated into broader hybrid warfare strategies and predicted that this trend would continue. Unsurprisingly, the past year's activities have reinforced that view. Geopolitical conflicts like the Russia-Ukraine war and the Israel-Gaza conflict continued to drive spikes in cyber activity, particularly across Europe, while state-sponsored APTs from China and Russia have maintained their focus on pre-positioning in Western critical infrastructure.

This chapter focuses on what's been changing in recent years – specifically, the erosion of boundaries between state-sponsored groups, cybercriminals and hackers, and the surge in ideologically motivated hacking. We also consider the structural factors that shape our perception of global cyber threats, and the intelligence gaps that these inherently create.

Key trends

- Pro-Russia hacktivist groups dominated the European threat landscape in 2025, accounting for the majority of incidents through high-volume but low-impact DDoS campaigns, closely tied to geopolitical flashpoints like the Russia-Ukraine war and Israel-Gaza conflict.
- The boundaries between state-sponsored APTs, cybercriminals, and hackers have eroded further as groups increasingly share tools and infrastructure, deliberately masquerade as one another, and mix espionage objectives with financially motivated operations.
- China, Russia, North Korea, and Iran maintained - and in some cases intensified - their cyber operations in 2025, with China focused on long-term pre-positioning in critical infrastructure, Russia targeting NATO allies, North Korea achieving record cryptocurrency theft, and Iran escalating attacks on Israel amid renewed regional conflict.

The rise of hacktivism

In 2025, hacktivism dominated the European threat landscape in terms of incident volume. According to ENISA, this ideologically motivated form of hacking was responsible for 79% of incidents targeting EU Member States, with DDoS attacks against government websites making up 91.5% of that total. But despite the volume, the actual impact remains limited: only 2% of hacktivist incidents resulted in meaningful service disruption.^[33]

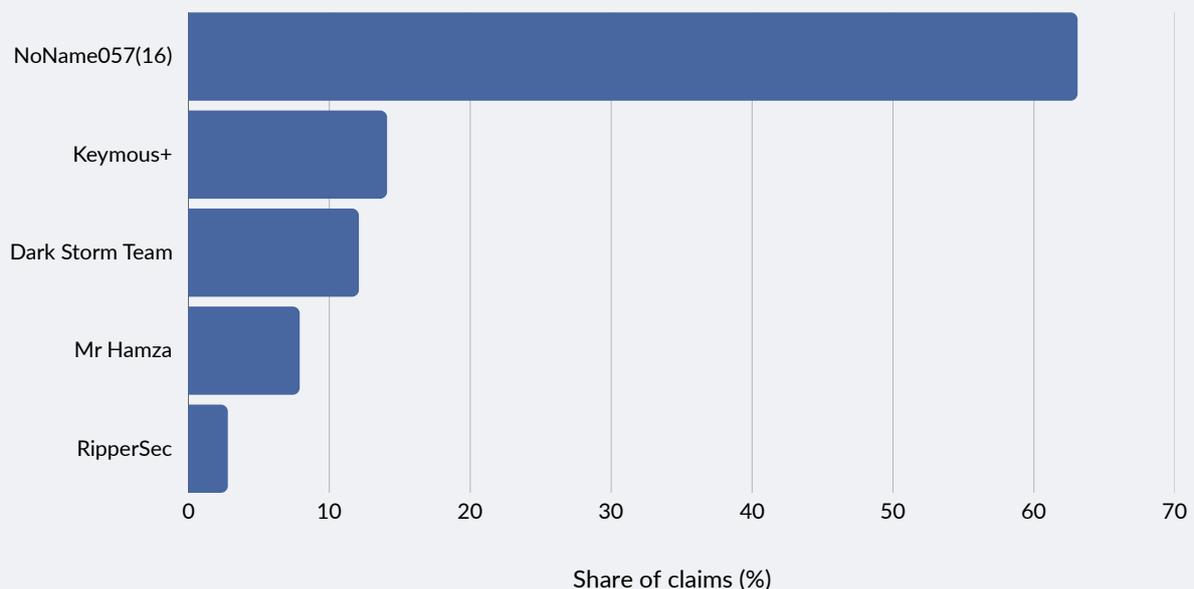
Pro-Russia hacktivism dominates this space. NoName057(16) alone claimed 63.1% of attacks, followed by Keymous+ (14.1%), Dark Storm Team (12.1%), Mr Hamza (7.9%), and RipperSec (2.8%).^[33]



Hacktivism is a blend of hacking and activism, where individuals or groups use digital tools to promote social or political causes - EBSCO

Of all these groups, NoName057(16) has sustained the highest operational tempo, running continuous campaigns throughout the reporting period.^[33] This consistency is likely enabled by their crowd-sourced model, operationalized through the DDoSia platform. DDoSia is a DDoS tool operated via Telegram, where NoName057(16) recruits volunteers (“heroes”) to install a client on their own devices and join coordinated attacks. Participants are rewarded based on their contributed attack traffic, effectively turning DDoS into a gamified, paid activity.^[34]

Top 5 groups claiming attacks against EU (2025)



Source: [ENISA](#)

NoName057(16) in the Netherlands

The Netherlands has been a frequent target of NoName057(16). In March 2025, a DDoS attack on government ICT administrator Logius temporarily took down DigiD, the national digital identity platform, blocking citizens from accessing services including the Tax Administration, UWV, DUO, and hospital patient portals. In April, 19 municipal and provincial websites were impacted by DDoS attacks and in June, during the NATO summit in The Hague, a dozen Dutch organizations, including several municipalities, faced coordinated attacks.^[35] A month later, law enforcement agencies from 12 countries, including the Netherlands, launched a joint operation against the group, resulting in the disruption of over 100 systems tied to the group's attack infrastructure and the takedown of several key servers.^[36]

Key drivers of hacktivism

The surge in hacktivist activities in Europe is closely linked to geopolitical events, namely the Russia-Ukraine war and Israel-Gaza conflict. Many attacks have been conducted against governments with direct involvement in the conflicts, or those who have announced support for either side. These actors are not subject to formal state control and tend to select targets based on vulnerability rather than strategic value, making their activities difficult to predict.

Although most of these campaigns are technically unsophisticated and cause little lasting disruption, they succeed in keeping political tensions in the headlines and eroding public confidence in digital services. That, for many of these groups, is the real objective. Looking ahead to 2026, we can expect hacktivism to remain a high-volume feature of the European threat landscape, with activity closely tracking geopolitical flashpoints.



Get the full overview of NoName057(16)

- Ideological motivations
- Recent campaigns
- Crowd-sourcing tactics
- Strategies to defend against them

[Read the blog](#)

Blurred lines

In last year's report, we observed that the boundaries between state-sponsored APTs, cybercriminals and hacktivists were becoming increasingly difficult to define. A year on, those lines have blurred even further.

After tracking this trend throughout 2025, it's clear that this convergence is being driven by three reinforcing factors: widespread sharing of tools and infrastructure, deliberate obfuscation of identities and affiliations, and "double dipping" by state-linked actors who mix espionage with financial crime.

Shared tools and infrastructure

Historically, intrusion sets could be distinguished by their level of sophistication, their objectives, or their operational tradecraft. In 2026, those markers are no longer reliable. Even advanced state-sponsored groups now routinely use commodity hacking tools, such as Malware-as-a-Service offerings, commercial proxy networks, and shared hosting infrastructure. This enables them to blend in with the noise of everyday cybercrime.

Russian groups APT29 and Sandworm, for instance, have been observed using commercial residential proxies and deploying commodity infostealers typically associated with financially motivated actors.^[37]

"APT and financially motivated ecosystems are increasingly overlapping, through dual-mission actors, shared tooling, and access brokerage. Highly capable actors often rely on commodity techniques because they're efficient and effective. As a result, distinguishing actor intent and sponsorship from TTPs alone has become increasingly difficult."

Jurjen Harskamp
CEO, Hunt & Hackett

Beyond using the same tools, we're seeing increasing collaboration between threat actors, particularly those from the same country. As Martijn Grooten notes, one cluster may focus on initial access and data theft, then sell that access on to ransomware operators if the target proves valuable. Even where there is no formal alliance, we see groups taking inspiration from successful campaigns of other actors.

For defenders, this convergence of tooling and TTPs makes traditional attribution based on technical indicators less reliable.

"We definitely see threat actors copying from each other. ClickFix is just one example from the past year."

Martijn Grooten
Threat Intelligence Analyst

Deliberate obfuscation

Some groups go further, actively masquerading as one another. An example of this is the trend dubbed *faketivism*, when state-aligned actors adopt hacktivist personas to obscure their true affiliation. Cyber Army of Russia Reborn, linked to Sandworm, and CyberAv3ngers, tied to Iran's Revolutionary Guard Corps, serve as two clear examples.^[38]

Cybercriminals, too, have been caught impersonating rival groups. One recent example involved scammers posing as CIOp ransomware group to extort businesses. The scammers claimed to have exfiltrated data from their targets' networks and threatened to post it on CIOp's blog if the ransom was not paid. However, researchers found the emails missing many of the signatures associated with genuine CIOp ransom demands, such as a 48-hour payment deadline and links to a secure chat channel for ransom payment negotiations.^[39]

Double dipping

Finally, state-sponsored actors are no longer confining themselves to strategic objectives. According to Verizon, approximately 28% of incidents involving state-affiliated groups in 2024 had a financial motive.^[40] In some cases, this is state-sanctioned - North Korea's cyber program, for instance, is explicitly designed to generate revenue for the regime. In others, it appears to be opportunistic rather than directed, with operators running financially motivated campaigns on the side to pad their own compensation.

The net effect is a threat landscape where motivations are less easily defined, attribution is more complex, and the traditional frameworks for understanding cyber threats are increasingly obsolete.

“Over the last years, we’ve learned much more about how these organizations operate, and we see a blend of criminal activity and state involvement. Stealing data, stealing secrets, gathering intelligence on important officials, or simply hitting organizations financially - all of it is intertwined in this space.”

Daan Weggemans

Assistant Professor, Leiden University

The Big Four

While hackers dominate the European threat landscape in terms of volume, state-sponsored actors from China, Russia, North Korea, and Iran continue to pose significant strategic risks. Each of these nations maintained or escalated their cyber operations in 2025, pursuing distinct objectives shaped by their unique geopolitical and economic conditions. The following profiles provide a snapshot of each country's capabilities, motivations, and notable campaigns from the past year.

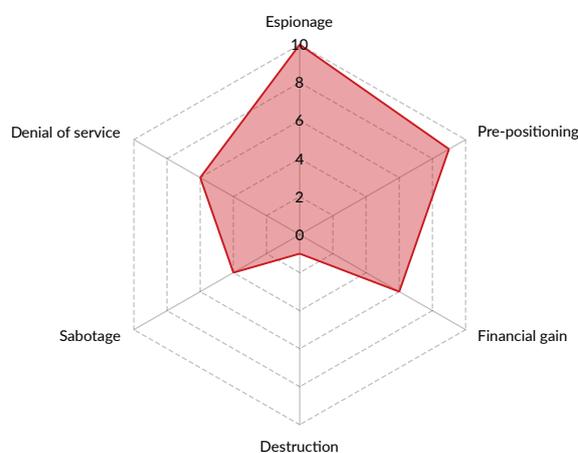




China

Strategic motivations

- Controlling dissent and competition through large-scale information collection.
- Establishing strategic deterrence by pre-positioning in foreign critical infrastructure.
- Gaining economic advantages through espionage and information theft.^[41]



2025 activities

China-nexus activity surged 150% across all sectors in 2025, reflecting continued investment in an extensive, mature cyber ecosystem that blends state intelligence services with universities, private companies and a volunteer cyber reserve.^[42] Operations remained focused on espionage and pre-positioning, with groups like Salt Typhoon and Volt Typhoon demonstrating remarkable adeptness at exploiting legitimate tools, credentials and processes to remain undetected in targeted networks for years at a time.

CASE STUDY

Typhoon Season

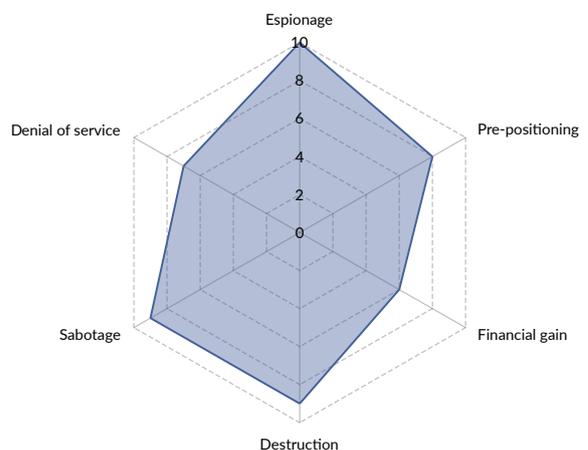
Over the past year, the full extent of Volt Typhoon's operations has come into sharper focus. The group has compromised IT environments across multiple US critical infrastructure sectors, including communications, energy, transportation, and water. In some cases, the group maintained access for up to five years. In one confirmed compromise of a water utility, Volt Typhoon actors gained access to PuTTY profiles for water treatment plants, wells, and an electrical substation, once again showcasing their ability to pivot from IT networks into operational technology (OT) systems.^[43]



Russia

Strategic motivations

- Gaining military and intelligence advantages in the context of the Ukraine war.
- Pre-positioning in Western critical infrastructure to enhance geopolitical leverage.
- Influencing political decisions and public narratives through disinformation campaigns.^[44]



2025 activities

In 2025, Russia's state-aligned cyber activity remained tied to the war in Ukraine, with attacks against NATO countries rising 25%. Nine of the 10 most-targeted nations were Alliance members - Ukraine was the tenth. Government bodies were hit hardest, followed by research, academic, and non-governmental organizations.^[45] Russia has maintained a high development tempo, deploying at least nine new wiper families and two ransomware variants since the beginning of the war.^[46] They are also experimenting with AI to enhance both malware capabilities and influence operations.^[47]

CASE STUDY

Laundry Bear

In May 2025, the Dutch intelligence service identified LAUNDRY BEAR, a new Russian state-backed threat actor linked to the September 2024 Dutch police breach. Active since 2024, the group conducts espionage against government, defense, and digital service providers across EU and NATO states, focusing on entities relevant to the Ukraine war. They specialize in rapidly compromising cloud email, exfiltrating mailboxes and Global Address Lists, and abusing valid accounts via pass-the-cookie attacks, password spraying, and living-off-the-land techniques.^[48]

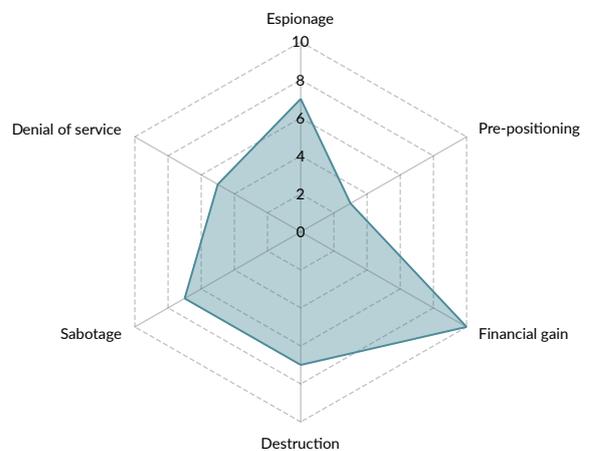


North Korea



Strategic motivations

- Funding the regime through large-scale cryptocurrency thefts.
- Compensating for limited domestic R&D through espionage campaigns.
- Gaining military and technological advantages.^[49]



2025 activities

North Korea continued its long-running "Operation Dream Job" espionage campaign, using recruitment-themed spear phishing to target European defense and aerospace firms.^[50] It also sustained a global effort to embed fraudulent North Korean IT workers into Western tech companies.^[51] It was also a record year for revenue generation, with DPRK-linked actors stealing an estimated \$2.02 billion in cryptocurrency, a 51% increase from 2024. This success appears to stem from a strategic pivot to targeting individual crypto holders, whose personal security tends to be weaker than that of major exchanges, and using embedded IT workers to gain privileged access for high-impact compromises.^[52]

CASE STUDY

ByBit hack

In February 2025, North Korea's Lazarus Group executed the largest cryptocurrency theft on record, stealing \$1.5 billion from exchange ByBit by compromising a supplier and redirecting 401,000 Ethereum tokens to attacker-controlled wallets. Within weeks, the group had converted at least \$300 million into unrecoverable funds. Analysts noted that the sophistication and speed with which the attackers were able to launder the stolen funds was indicative of a highly organized, dedicated operation.^[53]

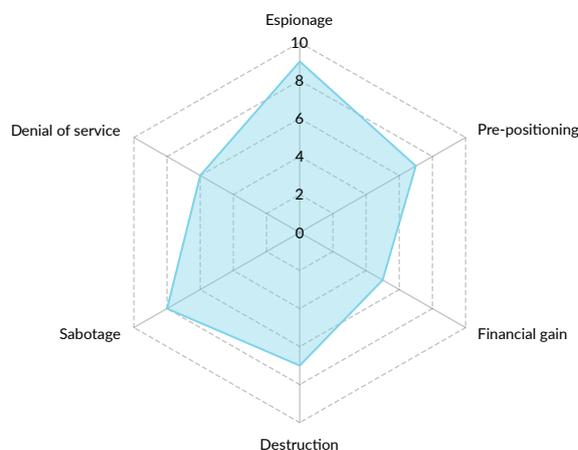


Iran



Strategic motivations

- Preserving the stability of the regime and projecting regional dominance.
- Asymmetric retaliation and deterrence against adversaries like the US and Israel.
- Countering sanctions and supporting domestic research and development.^[54]



2025 activities

Much of Iran's cyber activity in 2025 was driven by its regional conflict with Israel. The "12-Day War", which began on 13 June after Israeli strikes on Iranian weapons production sites, triggered a spike in cyberattacks from Iranian groups against Israeli government, defense, aerospace, and education targets.^[55] Among the most active was MuddyWater (affiliated with the Ministry of Intelligence and Security), which deployed a previously undocumented backdoor called MuddyViper against Israeli entities across multiple critical sectors. The campaign showed signs of growing operational maturity, including new tooling for stealth, persistence, and credential harvesting, as well as evidence of coordination with fellow Iran-aligned group Lyceum.^[56]

CASE STUDY

Charming Kitten leak

In November 2025, a significant leak of internal documents exposed the inner workings of Charming Kitten (also known as APT35 or Department 40), a cyber unit within Iran's Revolutionary Guard Corps known for targeting journalists, government agencies, and critical infrastructure across the Middle East, Europe, and North America. The leak revealed a direct connection between the group's cyber operations and targeted killings, showing that stolen information was fed into a system designed to locate and assassinate enemies of the Iranian regime.^[57]

Exclusive threat insights

FROM HUNT & HACKETT EXPERTS



Cozy Bear (APT 29)



Fancy Bear (APT 28)



Lazarus (APT38)



Double Dragon (APT41)



Sandworm



OilRig (APT34)



Charming Kitten



Silent Librarian



Sea Turtle

[Join the Members' Portal](#)

We're not seeing everything

The final "trend" in this section is really more of a health warning: our view of the global threat landscape is incomplete and rather biased. Cyber intrusions are covert by nature. Espionage actors invest heavily in remaining undetected, and many compromises are never discovered, let alone publicly disclosed. The incidents that do make headlines represent only a fraction of what's actually happening - and that fraction is often shaped by who has the incentive to talk about them.

"Countries always spy on each other and have always done that — and that includes allies. What we know based on public reporting is the tip of an iceberg, but it's a very biased tip."

Martijn Grooten
Threat Intelligence Analyst

When reading threat intelligence reports (including this one) it's worth asking: why is this incident being publicly reported? Who benefits from this narrative? Was disclosure mandated by regulation, driven by a victim's need to manage reputation, or used to support a particular geopolitical narrative? The answers often reveal as much as the content itself.

In the West, we tend to focus on the Big Four because they are our digital adversaries, and because criticizing them carries little political risk. Indeed, a long-running dataset tracking publicly known cyber operations since 2005 shows that China, Russia, Iran, and North Korea together account for roughly 77% of all suspected state-sponsored operations recorded globally.^[58]

This figure reflects attribution where researchers have sufficient evidence to link an operation to a nation - not the full picture of global cyber espionage. The uncomfortable truth is that everyone conducts espionage to varying degrees. Nations have spied on one another for centuries, first through conventional means, now predominantly via cyber intrusions. We just don't have the same level of visibility or talk much about allies targeting allies.

"There are countries that are safe to report on, such as Russia, Iran, North Korea, China - you're not going to upset anyone by calling them out. But when it comes to Western APT groups, there's little incentive to write about them."

Martijn Grooten
Threat Intelligence Analyst

Part of this visibility gap is also operational. Western intelligence agencies often have access to infrastructure-level capabilities, such as tapping undersea internet cables, that don't require targeting individual organizations unless there is a specific operational need.

This type of collection operates outside the organizational boundaries where most security monitoring takes place, leaving few traces in the telemetry that threat intelligence typically relies on. The 2013 Belgacom breach, in which a national telecommunications provider was compromised to enable mass surveillance, stands as a rare public example of this approach.

Recent geopolitical shifts have made this dynamic harder to ignore. Traditional alliances are showing signs of strain, and with that comes a growing awareness that threats don't only originate from familiar adversaries.

"We've always been more or less aware of the risks posed by Russia, China, Iran, and North Korea in the digital realm. But with recent political developments in the US, it's becoming more apparent that tensions also exist between countries that have traditionally been allies."

Daan Weggemans
Assistant Professor, Leiden University

None of this is to suggest that reporting on the Big Four or other adversarial nations is wrong. They sustain a high operational tempo, their activities are well documented, and they pose genuine risks. But a complete picture of the threat landscape would include actors we rarely discuss, conducting operations we seldom observe. That picture doesn't currently exist. The best we can do is remain aware of our intelligence gaps, study and experiment with attack techniques and read every report, including this one, with that caveat in mind.



CHAPTER FOUR

The Big Picture

Having looked at what attackers are doing and how geopolitics shapes the threat landscape, this final chapter turns to the foundations the cybersecurity industry is built on. We explore how Europe's growing preoccupation with digital sovereignty intersects with deep dependencies on foreign technology and threat intelligence, before ending with a frank assessment of the cybersecurity market itself - including the structural incentives and constraints that make meaningful progress harder than it should be.

Key trends

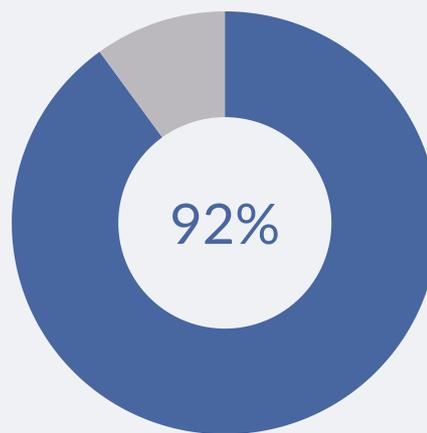
- Geopolitical shifts have placed Europe's deep reliance on American technologies and cloud infrastructure in the spotlight, intensifying calls for digital sovereignty. Experts warn that this is a complex ambition requiring pragmatic approaches to control technology, data, and regulatory alignment.
- Practitioners highlight that security dependence can quickly become geopolitical dependence. Reliance on non-EU security platforms and embedded threat intelligence has the potential to create structural detection gaps as priorities and alliances shift. Experts warn that if critical security capabilities depend entirely on non-EU technology, that dependence can be leveraged (directly or indirectly) through legal, commercial, or geopolitical pressure.
- The cybersecurity market faces several acute structural challenges that undercut meaningful improvements in collective resilience. This reality requires a shift in focus from tool accumulation to operational capability - controlling telemetry, validating detection logic, and treating compliance as a baseline rather than a finish line.

Digital sovereignty

Digital sovereignty became an increasingly popular topic in 2025, driven by rising geopolitical tensions that have forced Europe to confront uncomfortable questions about its dependence on American technology. The first shock came in March, when US President Trump's administration announced a pause on offensive cyber operations against Russia, representing a dramatic departure from previous US intelligence assessments that had consistently identified Moscow as a leading global threat.^[59] Around the same time, a US State Department official gave a speech at the United Nations about state-backed cyber threats, specifically mentioning China and Iran while omitting any mention of Russia.^[60] The signal to Europe was clear: when it comes to Russian intelligence, you're on your own. Many in the cybersecurity community feared the policy shift would embolden Russian actors further, removing a key deterrent and creating a clearer path for attacks against European targets.

Against this backdrop, concerns have increasingly been raised about Europe's overwhelming reliance on American cloud infrastructure. According to TNO, 92% of Western data is now hosted in the United States.^[61] And according to a recent NOS investigation, 67% of Dutch public sector and vital infrastructure domains rely on at least one US cloud service.^[62]

Given the current political climate, concerns have been raised about the ability of the US to access and potentially restrict European data. This would be possible under the US Cloud Act, which grants the government access to all data on American cloud services, even if the servers are located in Europe. This would naturally pose severe consequences for European organizations and individuals alike.^[63]



Share of Western data hosted in the United States.

Source: [TNO](#)

Similarly, heavy reliance on US threat intelligence introduces the risk of information gaps. Many European organizations depend on non-EU security platforms and embedded threat intelligence feeds. This dependence can create structural blind spots when priorities, legal frameworks, or geopolitical alignments shift - particularly if European organizations lack control over the underlying data, detection logic, or update governance.

It is not inconceivable, for instance, that US Cyber Command's "hunt forward" operations could gradually shift focus from Europe towards countries in China's sphere of influence, reducing the flow of intelligence on Russian cyber activities to European partners. Such a shift could also influence the commercial threat intelligence market, given that US government agencies are significant customers for many vendors. If demand for Russia-focused intelligence declines in the US, this may affect the availability and depth of such intelligence for European organizations who continue to face threats from Russia-nexus actors.^[64]

Taking back control?

These developments have prompted a lot of handwringing in Europe, followed by assertions that we should be taking back control of our data and digital technologies. At its core, this is what digital sovereignty means.

"Digital sovereignty is about technology, economy, and law. In order to be digitally sovereign as a country or on EU level, you need to be able to have control on the technology level, gain economic benefit from it, and ensure it adheres to your jurisdiction's highest legal standards."

Ellen Mok
Founder, Digitale Doetank

Most people agree this is an important ambition. Diving into the details, however, reveals enormous levels of complexity. Public discourse tends to focus on absolutes, largely ignoring the nuanced realities of global technological dependencies. Almost resoundingly, the experts we interviewed agreed that achieving complete, end-to-end sovereignty - whereby every layer of hardware, software, and hyperscale cloud infrastructure is Europe-owned and controlled - is simply not feasible in the current landscape.

"We need to move beyond binary thinking, full sovereignty or none at all. Instead, we should define which control points truly matter: what we need to own, control, or independently verify to maintain visibility, manage risk, and understand what's happening inside the services we depend on."

Jurjen Harskamp
CEO, Hunt & Hackett

Daan Weggemans echoes this, highlighting that even if desired, developing genuinely sovereign alternatives is "very expensive, it takes a long time and there's such a strong dependence that at some point it's nearly impossible."

The challenge stems from two major factors. First, while there is increasing interest in European alternatives, these cannot directly compete with the "one-stop shop" offerings of hyperscalers like Amazon, Microsoft and Google.

"If we wanted to move our SIEM or core hosting to a fully European stack, we'd have to accept real trade-offs in functionality and capacity. This isn't a short-term choice - it's a dependency built up over decades."

Ronald Prins
Co-founder, Hunt & Hackett

Decisions on where to accept such trade-offs will differ between organisations, particularly as a growing pool of European providers offers smaller-scale and often more niche solutions that may be suitable for specific needs.

Second, the dependencies embedded in modern technology stacks makes obtaining true digital independence, right down to the foundational layers of technology, exceptionally difficult to achieve in practice.

These dependencies extend far beyond the application layer, through the open-source libraries and licensing ecosystems that underpin modern software, and into the embedded software and hardware components whose production and maintenance are concentrated among a small number of global manufacturers. The risks created by these layered dependencies are visualised on the following page.

Dimensions of digital sovereignty

The digital sovereignty debate often centers on two high-level questions: who provides our technology, and where our data is stored. While important, this framing can flatten a far more complex reality. Modern technology stacks are layered and deeply interconnected, spanning cloud platforms, identity systems, open-source dependencies, embedded software, and hardware supply chains. Meaningful strategic choices require looking beneath the surface and understanding where control truly resides, where dependencies accumulate, and which layers introduce systemic risk. Only then can organisations define what must be owned, controlled, verified, or made substitutable.

The surface level

Public discourse tends to focus on vendor selection and data residency. While important, this framing overlooks the complex web of dependencies that underpin these services.

The software foundation

Modern software, regardless of where the primary vendor sits, is increasingly assembled from globally sourced open-source components and frameworks. These dependencies introduce two categories of risk. First, security risk: a single widely used vulnerability can propagate quickly across many products and environments. Second, operational and legal risk: maintainers can change licenses, revoke permissions, or shift terms in ways that force sudden rewrites, constrain usage, or disrupt critical functionality for downstream users.

The silicon core

At the deepest layer sits the firmware and embedded software that governs hardware behaviour, opaque by nature and difficult to independently validate. Vendor-controlled update channels and licensing mechanisms, combined with a highly concentrated hardware supply chain, mean geopolitical or commercial shocks can quickly translate into real operational risk.



“Many ‘European’ alternatives still rely on non-EU hardware and, by extension, firmware and embedded software, preserving supply-chain dependencies that are often overlooked in sovereignty debates. Acknowledging that reality enables a more practical discussion: what we need to observe, what we can control and independently verify, and what our true level of visibility and risk actually is.”

Jurjen Harskamp
CEO, Hunt & Hackett

None of this is to suggest that organisations should abandon the pursuit of digital sovereignty. Rather, it highlights the need to move beyond binary thinking and towards a more pragmatic, risk-based approach.

As Ellen Mok notes: “If you truly understand your stack and your dependencies, you can make conscious choices. Identify what is critical for your company and what should remain under European control. Then assess what kinds of data you handle, especially personal identifiers and sensitive information, that must not end up in the hands of other governments.”

The regulation paradox

Regulation is also a critical part of the equation, particularly in how it affects the global competitiveness of European technology providers. European privacy and data-protection frameworks place meaningful constraints on how AI systems can be trained, especially when it comes to the use of personal or sensitive data. These constraints reflect important societal values, but they also shape what is technically and commercially feasible.

In practice, US-based companies often operate in a regulatory and market environment that is less restrictive in certain areas (notably data use, product deployment and liability), which can enable them to aggregate larger datasets, iterate faster, and scale products more aggressively. This advantage is amplified by the structure of US venture and capital markets, which are designed to fund high-risk, high-reward outcomes and to create global category leaders. In Europe, funding is typically smaller, more fragmented, and often oriented toward de-risking and supporting incremental scale rather than winner-takes-most expansion.

These structural differences are visible in the AI race. The US benefits from a concentration of frontier labs, hyperscaler compute, and deep pools of capital creating compounding advantages in capability, scale, and adoption. Europe remains strong in research and regulation-led trust frameworks, but without comparable capital formation and scaling mechanisms, it risks becoming a downstream adopter rather than a global AI platform leader.

When procurement decisions are ultimately made by enterprises and European governments alike, the choice often falls on the solution that performs best at scale, even if it was developed under a regulatory baseline that would not be permissible within Europe. In practice, Europe risks exporting demand while importing dependency.

Jurjen Harskamp notes: "The result of this is a structural tension, where Europe enforces high standards at home but lacks equivalent mechanisms to ensure that those standards translate into competitive advantage for European firms, or into symmetrical obligations for external providers operating in the European market."

"We don't have a regulation problem, but we need to buy the products that have been designed with these regulations in mind. Otherwise, we just use regulations to push our own tech companies out of the market. And that's insane."

Ellen Mok
Founder, Digitale Doetank

"Without a meaningful stake in semiconductors, the cloud, AI, and cybersecurity, our long-term prosperity is at risk. It really is that simple."

Jurjen Harskamp
CEO, Hunt & Hackett

The digital sovereignty debate will only intensify as geopolitical tensions rise and technological dependencies deepen. Europe won't achieve complete independence overnight, but it can start making strategic choices about what to control and where to invest. The alternative means accepting both the risks and the long-term economic consequences of permanent dependence.

Tech as economic prosperity

The final aspect of this discussion relates to economic prosperity. Several of our experts noted that Europe's lack of a competitive domestic tech sector doesn't just leave us dependent on foreign infrastructure today, it undermines our chance of having a strong economy in the future.

As Ellen Mok notes: "Right now, our governments are heavily dependent on American technology. We spend billions in public funds every year effectively boosting a foreign economy - one that won't help us pay, for example, our pensions. So, we're investing taxpayers' money into an ecosystem that doesn't meaningfully contribute back to our own long-term wealth. That's a big problem."

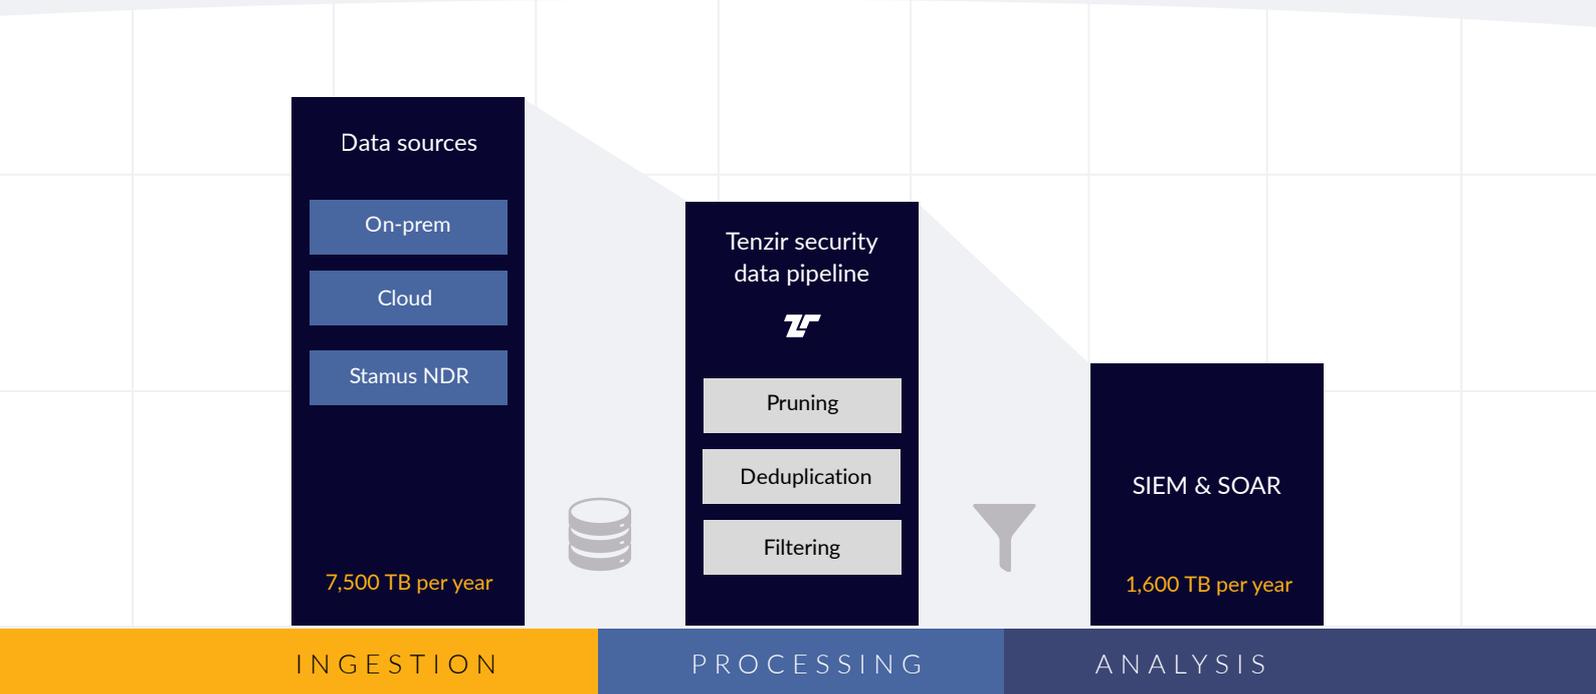
Achieving independent verifiability of security data

To increase autonomy over our security telemetry, Hunt & Hackett has built data pipelines that allow us to shape, enrich and route security data, independent of the security tools in which it is ultimately processed and stored. We develop these pipelines on Tenzir - an EU-based provider of a security-native data pipeline platform - which gives us control across the full chain, from collection and transformation to routing and action.

In practice, this has materially improved how we handle telemetry ingestion. For most data sources, we achieve 60–90% data reduction without loss of detection. We also maintain investigation-relevant visibility while producing a cleaner and more consistent dataset. The result is a significantly improved signal-to-noise ratio: fewer irrelevant fields, less ingestion waste, and higher-quality inputs for context, detection, threat hunting, and investigation.

Strategically, this architecture enables independent verifiability. Because we can model, enforce, and validate what is collected and retained before it reaches our SIEM/SOAR - for example by applying field-level allowlists, schema validation, and completeness checks - we can reason about coverage and completeness on our own terms rather than relying on opaque vendor processing logic. At the same time, it reduces vendor lock-in, lowers cost pressure driven by SIEM and telemetry economics, and turns raw logs into higher-value, analysis-ready data.

This is what our pragmatic approach to vendor-agnostic security and verifiability looks like in practice. It is also a concrete step toward sovereignty: not by attempting to replace every foreign dependency at once, but by identifying the first control points that matter - telemetry, transformation logic, and validation - and progressively building the capability to own, verify, and defend them. From there, we create the option to selectively shift additional capabilities (e.g., detections, hunting workflows) closer to the data layer, and to evolve hosting choices over time where it meaningfully reduces concentration risk. That includes improved portability across platforms and, over time, greater cloud/provider flexibility.



The cybersecurity market

Having traversed the threat landscape, from what we're seeing on the ground in the Netherlands (Chapter 1), to the evolving attacker playbook (Chapter 2), to cyber at the geopolitical frontlines (Chapter 3), and finally the broader structural forces shaping the market (Chapter 4), it feels only right to conclude by turning the lens inward. We are, after all, just another player in a crowded security marketplace, one with its own blind spots and incentives.

Every year, reports like this land in your inbox. Some are written in the spirit of information sharing. Others, candidly, are part of a market competing for attention and budget. Too often, the underlying message is the same: the threat landscape is worse, attackers are stealthier, and whatever you're doing still isn't enough. But the more important question is why it remains so difficult to defend against modern threats. Where are the structural gaps and systemic trade-offs that quietly erode security posture over time?

From our vantage point, three patterns keep repeating. First, much of the industry remains product-centric rather than threat-centric - more tools, more dashboards, more "visibility," yet teams feel no more in control. Second, many MDR offerings still default to passivity, waiting for alerts rather than continuously hunting for what static monitoring misses. Against high-stealth, living-off-the-land intrusions, that approach leaves organizations operating with an incomplete picture. Third, SIEM economics force a depressing choice between data completeness and cost. When telemetry is priced like a utility meter, organizations are incentivized to ingest less, store less, and investigate with fewer facts, just when the attacker playbook is accelerating in speed and scale.

Regulation adds another layer. Frameworks like NIS2 and DORA rightly push organizations toward resilience and provability. But when compliance becomes a checkbox exercise, it undermines the intent - security that looks good on paper yet fails under pressure.

This is also where digital sovereignty becomes practical, not ideological. Sovereignty is not "European or not," nor "cloud or not." It is about defining the control points that matter: who controls the telemetry, the detection logic, and the intelligence that shapes what gets flagged, and what gets missed. If we cannot independently validate data completeness, reproduce detections, and verify what our security stack is actually doing and observing, we inherit blind spots as vendor priorities, market incentives, and geopolitical alignments shift.

That is why our own direction of development is increasingly focused on control over telemetry and detection outcomes: designing data pipelines that reduce cost without sacrificing visibility; ensuring we can retain and query the right data for long enough to investigate; building detections and threat hunts that are portable and auditable; and continuously validating our assumptions through automated testing (Breach & Attack Simulation), so we can prove that the telemetry, detection logic and response processes actually fire as intended. We treat intelligence as something that must be testable in the environment, not just trusted because it came from a feed. In other words, resilience through transparency, verification, and the ability to adapt rather than reliance on assumptions.

Enterprise-grade security is about operational capability rather than adding more technology. It means retaining the right telemetry long enough to investigate, validating that it is complete, hunting proactively, and being able to demonstrate coverage and response in a way that holds up in the real world.

It also means repeatability and discipline at scale - consistent onboarding across environments, measurable detection quality, clear ownership and escalation paths, and the ability to operate under pressure, during incidents, audits and organizational change, without the security posture quietly degrading. Enterprise-grade is, ultimately, the difference between having tools and being able to reliably use them to prevent, detect, and contain intrusions.

Security is messy and imperfect. Pretending the path forward is simple does a disservice to the effort required. But by acknowledging the trade-offs, designing for verifiability, and treating compliance as a baseline, as opposed to a finish line, organizations can take meaningful steps toward genuine resilience.

So what should I do on Monday?

If there is one takeaway from this report, it's that resilience doesn't come from adding more tools. It comes from understanding where you're exposed, what you actually see, and whether your controls work in practice.

Concretely, moving toward enterprise-grade and sovereign-ready security starts with:

- Map your real control points** – identity systems, edge devices, cloud control planes, CI/CD pipelines, and third-party integrations – and be honest about what's monitored versus left blind.
- Validate identity assumptions** – would credential, token, service account, or helpdesk abuse look “normal” in your logs, and would anyone notice?
- Reduce blast radius** – tighten privileged access (least privilege, tiered admin, JIT/JEA, break-glass discipline) so one identity doesn't become total compromise.
- Know your telemetry** – what you ingest, how long you retain it, and what you've excluded because of cost or complexity. Assume attackers operate in those gaps.
- Test detections, don't trust them** – validate high-impact attack paths (identity abuse, edge compromise, cloud lateral movement, data staging) and prove alerts actually fire.
- Treat edge and cloud exposure as first-class risks** – not hygiene issues, but primary initial access vectors.
- Look at your supply chain through an attacker's lens** – where trust is inherited, compromise scales (including software supply chain and code-signing trust).
- Make recovery a first-class control** – prove you can restore and rebuild the identity/control plane under pressure, not just “detect faster.”
- Define sovereignty as control** – what telemetry, detection logic, and intelligence you must independently verify, adapt, and retain if vendor priorities or geopolitics shift.
- Use compliance as a baseline, not a finish line** – evidence resilience, don't just document it.

Enterprise-grade security is not more technology. It is operational capability, verifiability, and control. This is the short version. For the complete checklist (with the full set of enterprise-grade and sovereignty-ready actions), see Appendix: “So what should I do on Monday?”

References

1. Eye Security. (2024, December 17). Eye Security warns: BEC incidents surge in 2024, driving up insurance costs. <https://www.eye.security/press/bec-incidents-surge-in-2024-driving-up-insurance-costs>
2. Aggarwal, V. (2022, August 24). Why business email compromise still tops ransomware for total losses. CSO Online. <https://www.csoonline.com/article/573435/why-business-email-compromise-still-tops-ransomware-for-total-losses.html>
3. Amr. (2025, November 18). Understanding technical debt: The new zero day. Bugcrowd. https://www.bugcrowd.com/blog/understanding-technical-debt-the-new-zero-day/?utm_campaign=blog&utm_source=linkedin&utm_medium=organic_social&utm_content=1764090045
4. Mucci, T. (2025, November 17). *Technical debt*. IBM. <https://www.ibm.com/think/topics/technical-debt>
5. Jackson, C. (2025, October 24). *Extortion without encryption: the next phase of ransomware*. CyberMaxx. <https://www.cybermaxx.com/resources/extortion-without-encryption-the-next-phase-of-ransomware/>
6. 2025 Data Breach Investigations Report. (2025). In *Verizon Business*. Retrieved February 17, 2026, from <https://www.verizon.com/business/resources/reports/dbir/>
7. *Amazon Threat Intelligence identifies Russian cyber threat group targeting Western critical infrastructure*. (2025, December 23). Amazon Web Services. <https://aws.amazon.com/blogs/security/amazon-threat-intelligence-identifies-russian-cyber-threat-group-targeting-western-critical-infrastructure/>
8. The Hacker News. (2025, August 28). *Salt typhoon exploits flaws in edge network devices to breach 600 organizations worldwide*. <https://thehackernews.com/2025/08/salt-typhoon-exploits-cisco-ivanti-palo.html>
9. *PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure*. (2024, February 7). CISA. Retrieved February 17, 2026, from <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>
10. CrowdStrike. (2025). 2025 Global Threat Report. <https://www.crowdstrike.com/en-us/global-threat-report/>
11. Valid Accounts, technique T1078 - Enterprise | MITRE ATT&CK. (n.d.). <https://attack.mitre.org/techniques/T1078/>
12. Cloud Protection for Salesforce by WithSecure. (2026, January 27). *Salesforce attacks in 2025: Why cyber criminals are targeting Salesforce*. <https://cloudprotection.com/blog/salesforce-attacks-in-2025/>
13. 2025 Data Breach Investigations Report. (2025). In *Verizon Business*. Retrieved February 17, 2026, from <https://www.verizon.com/business/resources/reports/dbir/>
14. Global Cybersecurity Outlook 2026. (2026). In World Economic Forum. <https://www.weforum.org/publications/global-cybersecurity-outlook-2026/digest/>
15. ENISA Threat Landscape 2025. (2025). In ENISA. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>
16. *IBM X-Force 2025 Threat Intelligence Index*. (2025, April 16). IBM. <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/2025-threat-intelligence-index>
17. Henderson, C. (2025, April 2). *X-Force Threat Intelligence Index 2024 reveals stolen credentials as top risk, with AI attacks on the horizon*. IBM. <https://www.ibm.com/think/x-force/2024-x-force-threat-intelligence-index>
18. *64% Infection Rate from Infostealer Malware in Businesses*. (2025, August 20). Hadrian. <https://hadrian.io/blog/64-infection-rate-from-infostealer-malware-in-businesses>
19. *IBM X-Force 2025 Threat Intelligence Index*. (2025c, April 16). IBM. <https://www.ibm.com/thought-leadership/institute-business-value/report/2025-threat-intelligence-index>
20. Investigation report on Jaguar Land Rover cyberattack - CYFIRMA. (2025, September 24). <https://www.cyfirma.com/research/investigation-report-on-jaguar-land-rover-cyberattack/>
21. *Clipboard to Compromise: PowerShell Script Self-Pwn*. (2024, July 29). Proofpoint. <https://www.proofpoint.com/us/blog/threat-insight/clipboard-compromise-powershell-self-pwn>

References

22. Coker, J. (2026, February 16). *ClickFix attacks surge 517% in 2025*. Infosecurity Magazine. <https://www.infosecurity-magazine.com/news/clickfix-attacks-surge-2025/#:~:text=ClickFix%20social%20engineering%20attacks%20have,the%20end%20of%20the%20year.>
23. Lakshmanan, R. (2025a, March 19). *ClearFake infects 9,300 sites, uses fake ReCAPTCHA and Turnstile to spread Info-Stealers*. The Hacker News. <https://thehackernews.com/2025/03/clearfake-infects-9300-sites-uses-fake.html>
24. *IBM X-Force 2025 Threat Intelligence Index*. (2025c, April 16). IBM. <https://www.ibm.com/thought-leadership/institute-business-value/report/2025-threat-intelligence-index>
25. Terry, R. (2025, March 26). *What is a Kerberoasting Attack?*. CrowdStrike. <https://www.crowdstrike.com/en-us/cybersecurity-101/cyberattacks/kerberoasting/>
26. Google Threat Intelligence Group. (2025, November 5). *GTIG AI Threat Tracker: Advances in threat actor usage of AI tools*. Google Cloud Blog. <https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools/>
27. CrowdStrike. (2025). *2025 Global Threat Report*. <https://www.crowdstrike.com/en-us/global-threat-report/>
28. Harme, J. (2025, October 30). *UK retailers hit by surge in AI-powered phishing & vishing scams*. Security Brief. <https://securitybrief.co.uk/story/uk-retailers-hit-by-surge-in-ai-powered-phishing-vishing-scams>
29. CrowdStrike. (2025). *2025 Global Threat Report*. <https://www.crowdstrike.com/en-us/global-threat-report/>
30. Google Threat Intelligence Group. (2025, November 5). *GTIG AI Threat Tracker: Advances in threat actor usage of AI tools*. Google Cloud Blog. <https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools/>
31. Anthropic. (2025). *Disrupting the first reported AI-orchestrated cyber espionage campaign*. In *Full Report*. anthropic.com. <https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf>
32. Stein, S. (2025, December 11). *CrowdStrike researchers identify hidden vulnerabilities in AI-Coded software*. CrowdStrike. <https://www.crowdstrike.com/en-us/blog/crowdstrike-researchers-identify-hidden-vulnerabilities-ai-coded-software/>
33. *ENISA Threat Landscape 2025*. (2025). In *ENISA*. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>
34. Insikt Group. (n.d.). *Inside DDOSIA: NoName057(16)'s Pro-Russian DDOS campaign infrastructure*. <https://www.recordedfuture.com/research/anatomy-of-ddosia>
35. Moester, M. D. & T. (2025, June 26). *Rising Cyber Tensions: NoName057(16)'s hacktivist activities reach the Netherlands*. Hunt & Hackett. <https://www.huntandhackett.com/blog/noname05716s-hacktivist-activities-reach-the-netherlands>
36. Europol. (2025, August 1). *Global operation targets NoName057(16) pro-Russian cybercrime network*. <https://www.europol.europa.eu/media-press/newsroom/news/global-operation-targets-noname05716-pro-russian-cybercrime-network>
37. CrowdStrike. (2025b, October 21). *2025 European Threat Landscape Report*. <https://www.crowdstrike.com/en-us/resources/reports/2025-european-threat-landscape-report/>
38. *ENISA Threat Landscape 2025*. (2025). In *ENISA*. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>
39. Coker, J. (2026b, February 16). *Fraudsters impersonate CLOP ransomware to extort businesses*. Infosecurity Magazine. <https://www.infosecurity-magazine.com/news/fraudsters-clop-ransomware-extort/>
40. *2025 Data Breach Investigations Report*. (2025). In *Verizon Business*. Retrieved February 17, 2026, from <https://www.verizon.com/business/resources/reports/dbir/>
41. *China's national cyberspace security strategy*. (n.d.). Digital Watch Observatory. <https://dig.watch/resource/chinas-national-cyberspace-security-strategy>
42. CrowdStrike. (2025). *2025 Global Threat Report*. <https://www.crowdstrike.com/en-us/global-threat-report/>

References

43. PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure. (2024, February 7). CISA. Retrieved February 17, 2026, from <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>
44. Jones, S. G. (2025). *Russia's shadow war against the West*. <https://www.csis.org/analysis/russias-shadow-war-against-west>
45. Milmo, D. (2025, October 16). *Russian cyber-attacks against Nato states up by 25% in a year, analysis finds*. The Guardian. <https://www.theguardian.com/world/2025/oct/16/russian-cyber-attacks-against-nato-states-up-by-25-in-a-year-analysis-finds>
46. Chincharadze, K. (2025, July 30). *From Georgia to Ukraine: Seventeen years of Russian cyber capabilities at war*. Modern War Institute. <https://mwi.westpoint.edu/from-georgia-to-ukraine-seventeen-years-of-russian-cyber-capabilities-at-war/>
47. Paganini, P. (2025, October 10). *Ukraine sees surge in AI-Powered cyberattacks by Russia-linked Threat Actors*. Security Affairs. <https://securityaffairs.com/183222/apt/ukraine-sees-surge-in-ai-powered-cyberattacks-by-russia-linked-threat-actors.html>
48. AIVD en MIVD onderkennen nieuwe Russische cyberactor. (2025, May 27). AIVD. <https://www.aivd.nl/documenten/publicaties/2025/05/27/aivd-en-mivd-onderkennen-nieuwe-russische-cyberactor>
49. Russell, D. (2019, April 30). *North Korea's Next Weapon of Choice: Cyber*. Asia Society Magazine. <https://asiasociety.org/magazine/article/north-koreas-next-weapon-choice-cyber>
50. Pichon, M., & Bonnefoi, A. (2025, November 20). *A Pain in the Mist: Navigating Operation DreamJob's arsenal*. Orange Cyberdefense. <https://www.orange cyberdefense.com/global/blog/cert-news/a-pain-in-the-mist-navigating-operation-dreamjobs-arsenal>
51. Sabin, S. (2025, August 19). *How North Korea's IT army is hacking the global job market*. Axios. <https://www.axios.com/2025/08/19/north-korea-it-worker-fraud-fortune-500>
52. Tidy, J. (2025, October 7). *North Korean hackers stealing record sums, researchers say*. <https://www.bbc.com/news/articles/cwy8z7wx03o>
53. Tidy, J. (2025a, March 10). *North Korean hackers cash out hundreds of millions from \$1.5bn ByBit hack*. <https://www.bbc.com/news/articles/c2kgndwwd7lo>
54. 2025 Data Breach Investigations Report. (2025). In *Verizon Business*. Retrieved February 17, 2026, from <https://www.verizon.com/business/resources/reports/dbir/>
55. *Israel and Iran on the brink: Preventing the next war*. (2025, October 3). European Union Institute for Security Studies. <https://www.iss.europa.eu/publications/briefs/israel-and-iran-brink-preventing-next-war>
56. Lakshmanan, R. (2025, December 2). *Iran-Linked Hackers Hit Israeli Sectors with New MuddyViper Backdoor in Targeted Attacks*. The Hacker News. <https://thehackernews.com/2025/12/iran-linked-hackers-hits-israeli-2.html>
57. Gharib, N. (2025, November 3). *Department 40 Exposed: Inside the IRGC Unit Connecting Cyber Ops to Assassinations*. <https://blog.narimangharib.com/posts/2025%2F11%2F1763938840948?lang=en>
58. *Cyber Operations Tracker*. (n.d.). <https://www.cfr.org/cyber-operations/>
59. FitzGerald, J. (2025, March 3). *Hegseth orders pause in offensive US cyber operations against Russia*. <https://www.bbc.com/news/articles/c2er34w0jgdo>
60. Jackson, J. (2025, March 17). *Trump is giving Russian cyber ops a free pass – and putting western democracy on the line*. TBIJ. <https://www.thebureauinvestigates.com/stories/2025-03-07/trump-is-giving-russian-cyber-ops-a-free-pass-and-putting-western-democracy-on-the-line>
61. *Test environment for cloud services: Structura-X*. (n.d.). TNO. <https://www.tno.nl/en/digital/digital-infrastructures/innovations-developments-european-cloud/>

References

62. Kieviet, E. (2026, January 27). *Nederlandse cloud "binnen handbereik", Tweede Kamer wil haast maken*. NOS.
<https://nos.nl/artikel/2599837-nederlandse-cloud-binnen-handbereik-tweede-kamer-wil-haast-maken>
63. José van Dijck: *'Europe has lost its digital sovereignty, and here's how to reclaim it'* - KNAW. (n.d.).
<https://www.knaw.nl/en/news/jose-van-dijck-europe-has-lost-its-digital-sovereignty-and-heres-how-reclaim-it>
64. Paulus, A. (n.d.). *Europe's cybersecurity depends on the United States*. Stiftung Wissenschaft Und Politik (SWP).
<https://www.swp-berlin.org/en/publication/europes-cybersecurity-depends-on-the-united-states>

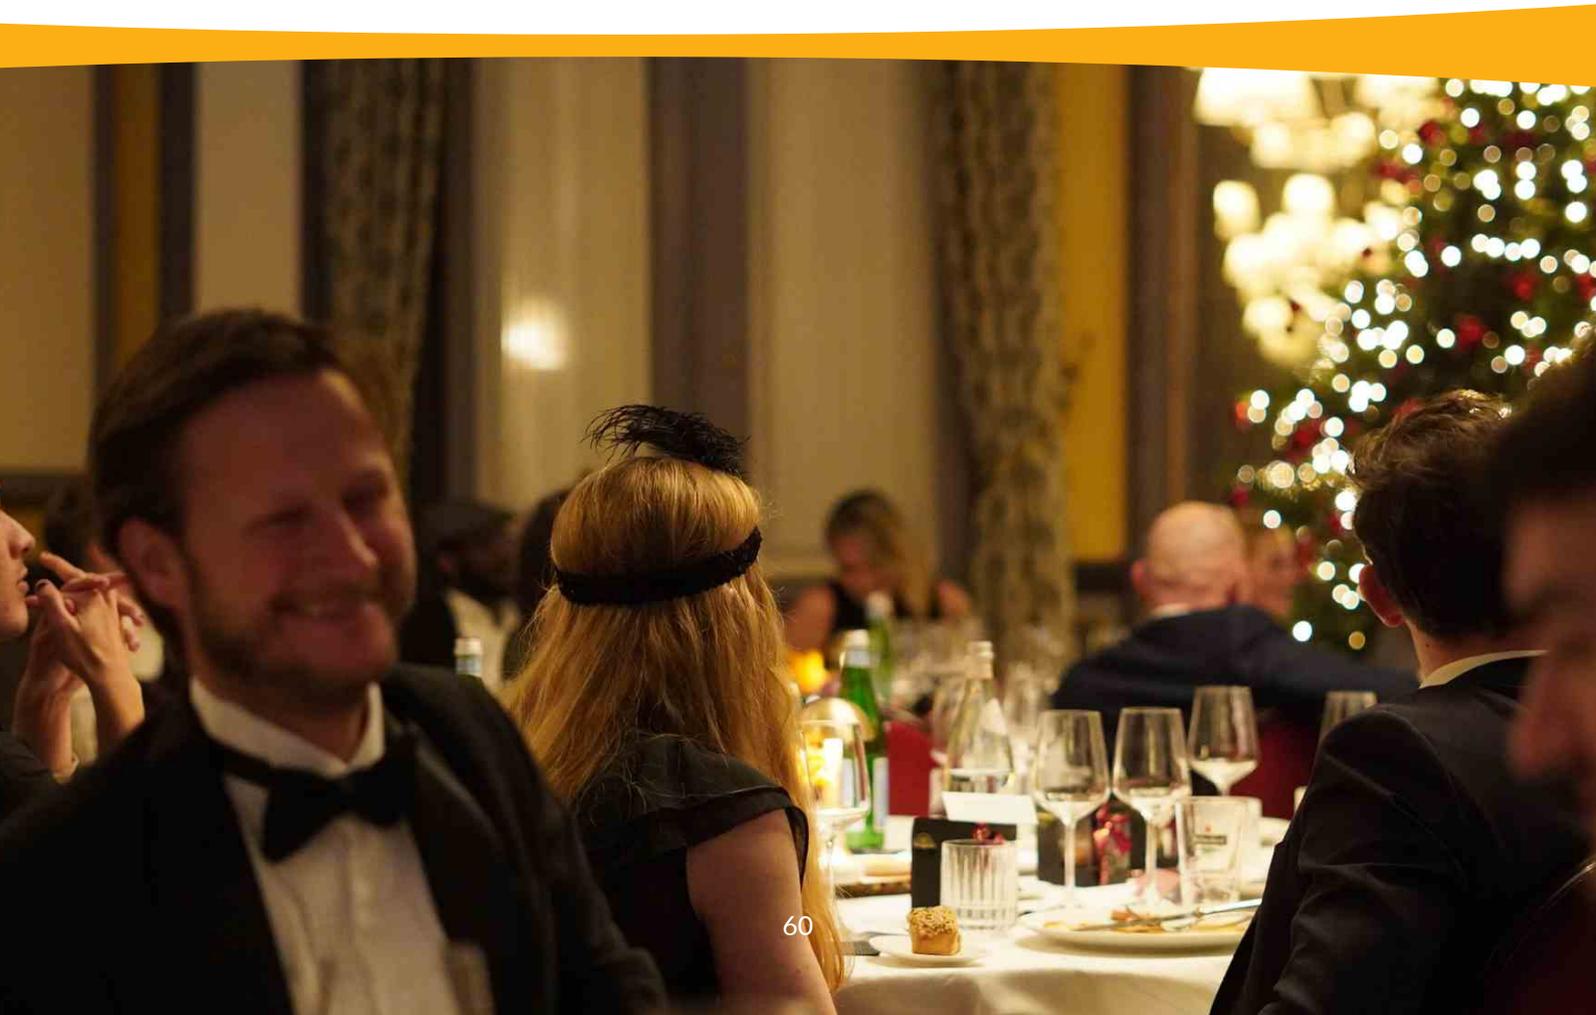
HUNT & HACKETT

Hunt & Hackett was founded in 2020 by Ronald Prins (co-founder Fox-IT) and Jurjen Harskamp (former executive Fox-IT) to help European companies prevent, detect and respond to today's most advanced adversaries. Leveraging threat modelling and data science, Hunt & Hackett builds, operates and maintains digital immune systems to protect against Advanced Persistent Threat (APT) groups and less sophisticated cybercrimes such as phishing and ransomware. Today, the company has grown to a team of around 60 people from a wide range of backgrounds and disciplines, including red teamers, threat intelligence experts, data scientists, and engineers. Outside of the office, we like to bring the team together to enjoy each other's company in unique locations.

Raison d'être

Ronald and Jurjen started Hunt & Hackett because they recognised that European knowhow, technology and intellectual property is increasingly being targeted by espionage campaigns. Compliance and technology-driven strategies are still the norm, yet these strategies are increasingly exposed by today's APT groups. This can cause unparalleled disruption and irreversible damage.

They also noticed that many CISOs struggle with developing the capabilities required to outsmart these digital adversaries. Why? Because cybersecurity readiness only gets added to the CxO agenda after a significant breach has occurred. Security partners, offering technology and process-based solutions, often lack the expertise and intelligence required to develop effective defences against these actors. Hunt & Hackett aims to bridge this gap, helping organisations in a wide range of sectors become more resilient to cyber threats.



Appendix

A practical checklist for moving toward enterprise-grade, sovereignty-ready security

This appendix translates the trends and structural challenges in this report into concrete actions you can take. It is not a maturity model or a shopping list, but it is built on a simple premise: resilience comes from what you can see, control, and verify.

The aim is to help organisations move toward enterprise-grade security that holds up under real attack, regulatory scrutiny, and organisational stress - while also taking a first step toward sovereignty-ready security by retaining meaningful control over telemetry, detection logic, and security decision-making, even as vendors, markets, or geopolitics shift.

You don't need to do everything at once. But you do need to know where you stand. These actions focus deliberately on visibility, control, and verification, the common failure points across identity abuse, cloud compromise, supply-chain attacks, and geopolitical pre-positioning described in this report.

Complete “So what should I do on Monday?” checklist

- Map your real control points**
 - Identity, edge devices, cloud control planes, CI/CD, and third-party/SaaS integrations - where attackers actually enter and move.
- Validate your identity assumptions**
 - Credentials, session tokens, service accounts, MFA resistance (AiTM), and helpdesk processes (incl. vishing). Ask: would abuse look “normal” in logs?
- Reduce blast radius via privileged access discipline**
 - Tiered admin model, least privilege, JIT/JEA, break-glass, secrets management/rotation, service account governance, strong segmentation.
- Check your telemetry coverage (and integrity), not just your tools**
 - What is ingested, normalized, and retained - and where can logs be bypassed/tampered with (cloud audit logs disabled, forwarding failures, time sync issues)?
- Fix the edge loop: exposure → patch/mitigate → verify**
 - Continuous external exposure monitoring, hardening baselines, patch/mitigation SLAs, and verification that fixes actually landed (esp. internet-facing appliances).
- Treat cloud as an identity and control-plane problem**
 - Inventory admin paths, OAuth consents, API tokens, conditional access, cloud audit logging, and “what happens if this admin account is abused?”
- Look at your supply chain through an attacker’s lens**
 - Where trust is inherited: OAuth permissions, shared admin, signed updates, vendor access, and downstream blast radius.

- Explicitly secure the software supply chain (CI/CD + code signing)**
 - Build pipelines, artifact registries, package managers, signing keys, build isolation, and provenance/attestation (prevent “trusted update as delivery channel”).
- Test, don't trust, your detections (and assumptions)**
 - Pick high-impact attack paths (identity abuse, edge compromise, cloud lateral movement, data staging) and validate detections with BAS / purple teaming.
- Hunt proactively where stealth lives**
 - Repeatable threat hunts for identity anomalies, token abuse, living-off-the-land, and long-dwell behaviour—backed by sufficient retention.
- Prepare for data extortion (not just encryption)**
 - Monitor for staging/exfil, impossible travel + token reuse, bulk mailbox pulls, unusual egress, and “name-and-shame” pressure paths.
- Make recovery a first-class control**
 - Immutable backups, restore tests, rebuild identity/control planes (AD/Azure AD, key SaaS), IR runbooks, and defined RTO/RPO that holds under real ransomware pressure.
- Operationalise governance for “enterprise-grade”**
 - Clear ownership, escalation paths, comms/legal/regulatory playbooks, and tabletop exercises for the scenarios you describe (identity takeover, edge compromise, SaaS token theft, supplier breach).
- Use compliance as a baseline, not a finish line**
 - Translate NIS2/DORA requirements into visibility, testing, and evidence that holds up during audits and incidents.
- Define “sovereignty” for your organisation as control points**
 - Not “EU vs non-EU”, but: what telemetry you must retain, what detection logic you must be able to inspect/extend, what you must independently verify, and how you can exit/substitute if priorities or geopolitics shift.
- Measure what matters (and publish it internally)**
 - Coverage of control points, telemetry completeness, detection quality (TP rate), time-to-contain, restore time, and the percentage of critical paths tested.
- Harden email + user interaction controls (AiTM, vishing, ClickFix)**
 - Resistant MFA methods, browser/session protections, conditional access/device posture, and service desk verification standards.
- Add posture metrics that prove progress**
 - MTTD/MTTR is fine, but also the percentage of critical control points fully logged and tested, plus detection quality (true positive rate, dwell-time assumptions).

