

Cyber Resilience Insights 2026

64%

of confirmed incidents
involved phishing, malware
or misused identity

43%

of confirmed incidents
began outside weekday
business hours

Every

forensic investigation
found exposed credentials,
tokens or secrets

27%

of rated offensive
security findings were
Critical or High



We see no quiet season in cybersecurity attacks; vigilance is necessary throughout the year.



Foreword from our CEO

This first edition of our cybersecurity report shows that **64%** of confirmed incidents involved phishing, malware or misused identity, and **43%** began outside weekday business hours.

Fast detection only delivers its full value when organisations can make and execute decisions just as quickly. The challenge for leaders, as I see it, is not to predict every threat that might emerge, but to strengthen the foundations that determine how effectively their organisations can respond, recover, and adapt.

Six important leadership decisions for 2027

- 01 Fund phishing-resistant identity**
– for people and machines.
- 02 Open a legacy-retirement lane** – removing unsupported technology and persistent exposure.
- 03 Treat recovery and management systems as Tier Zero** – protecting the systems needed to manage and restore the business.
- 04 Govern software supply chain and pipeline trust** – controlling the access and authority of trusted automation.
- 05 Authorise containment outside business hours** – so response does not depend on waiting for the working day to begin.
- 06 Make forensic readiness measurable** – ensuring organisations can establish what happened, how far an incident reached and what risk remains.

Has AI changed the game?

AI presents the next major challenge for cyber resilience. While it may not make the fundamentals of cybersecurity obsolete, it will change the pace and scale of cyber risk. At the same time, deploying AI agents as trusted actors with access to real systems, data and workflows makes trust itself a new attack surface. SOC's must therefore prepare to monitor, govern and contain AI-enabled activity with the same urgency they apply to users, machines and privileged access; strong identity, reduced attack paths, protected management environments and rapid containment remain essential.

Where do the conclusions come from?

At Conscia, we see the reality of cyber risk every day through our Security Operations Centres (SOC) in Slovenia, Denmark and Spain.

This report contains evidence from Conscia's SOC and Managed Detection & Response operations and is based on **260,943 investigations** performed between **1 July 2025** and **11 June 2026**, of which **1,266 were confirmed as security incidents and successfully contained.** It then combines these insights with digital forensics and offensive security testing.

I would like to thank everyone who contributed their expertise and experience to this report, especially its author, **David Kasabji**, for bringing together the analysis that follows.

Happy reading

Lars Faeste
CEO, Conscia

Table of contents

Foreword from our CEO.....	2	03 Digital forensics and incident response	17
01 Introduction.....	5	When intrusions succeed, evidence decides.....	18
How to read this report	6	Different incidents, same trust problem.....	19
Executive brief	7	04 Offensive security	21
Executive agenda.....	8	What controlled testing proved	22
Executive scorecard	10	Four failures that keep returning	23
Evidence.....	12	The treatment model	24
02 Managed Detection & Response.....	12	05 The recurring threats across our evidence	25
What reached the SOC, and what mattered	13	Where the evidence points in the same direction.....	26
There was no quiet season.....	14	06 The wider picture.....	27
How attacks really arrive.....	15	What the industry saw, and what we saw.....	28
Detection is fast. The clock is faster.....	16	Three insights from the industry, read against our year	29
		07 Outlook.....	31
		AI did not change the attack. It changed the pace.....	32
		AI agents will turn trust into the next attack surface	33
		The risks already taking shape for 2027	34
		08 Appendix	36
		Method, scope and limitations	36
		Glossary.....	37
		About the report's author	37



”

Organisations do not need to predict every threat that will emerge in 2027. They need to strengthen the foundations that determine how effectively they can prevent attacks, respond and recover.

Lars Faeste
CEO, Conscia



01

Introduction

How to read this report

Evidence, not opinion

All findings and insights presented in this report are derived from activities undertaken by our teams during the reporting period (1 July 2025 – 11 June 2026). We are committed to an evidence-based approach – where the available data does not support a conclusion, we do not make one, and where uncertainty exists, it is clearly acknowledged.

Three sources

Monitoring, investigation, and testing

Monitoring shows what happened daily, whereas forensic investigation demonstrates what happened when defences failed and offensive testing points to what could have happened. Evidence can come from one source, but confidence in it increases when sources converge.

A fourth lens

Read alongside the global picture

After the initial analysis, we read our conclusions alongside the year's major public threat research. Chapter 5 sets the two side-by-side, so that our observations can be interpreted in the context of findings from the wider industry.

Confidence language

Confirmed, assessed, but not determined

Forensic conclusions are assigned a confidence level. Confirmed findings are supported by direct evidence, while assessed findings are supported by indirect evidence and analytical judgement. Where the root cause or full scope cannot be determined, we state this explicitly rather than speculate.

What we withhold

Anything that identifies a client

Client identities, client counts, case counts and engagement counts are all withheld. Where a denominator would be revealing we report the proportion or the frequency band, but withhold the denominator itself.

ⓘ A note on comparison

This is the first edition of our Cyber Resilience Insights report and serves as a baseline for future analysis. As a result, we do not present year-on-year trends or comparisons.

Executive brief

The risk landscape at a glance

The year's central finding was not that attackers became dramatically more novel. It was that four familiar conditions repeatedly combined into serious outcomes – and each appeared across all three of our evidence sources (monitoring, investigation, and testing).

01 Identity

Credentials prove to be the common denominator

Phishing, malicious links and misuse of valid accounts were the most common confirmed incident types. But the deeper pattern was identity exposure: every forensic investigation in the supplied investigation set involved exposed credentials, tokens or secrets, and offensive testing showed similar paths through weak passwords and stolen sessions. In short, attackers succeed when trusted access can be stolen, reused, or left active for too long.

02 Trust

Legitimate workflows become attack paths

Remote access, APIs, cached tokens, scripting and build automation allowed malicious activity to blend into normal operations. In some environments, administrative and recovery systems were also reachable from ordinary user segments. Trusted workflows became attack paths.

03 Visibility

The business risk posed by uncertainty

Every investigation encountered material evidence gaps, including missing perimeter logs, interrupted telemetry or incomplete runtime visibility. Organisations cannot confidently define the scope of an incident when evidence is incomplete, so recovery decisions are made on assumptions rather than facts.

04 Time

Attackers do not wait for office hours

More than four in ten confirmed incidents began at night or during the weekend. Detection was fast in most confirmed cases, but the ability to authorise action did not always move at the same speed.

➔ What this means for executives

These findings point to four executive-level questions that amount to more than just technical observations. Are critical identities protected and revocable? Are trusted systems and workflows governed as potential attack paths? Can the organisation prove what happened when an incident occurs? And can containment be authorised when the incident clock starts, not when office hours resume? The next page turns these questions into six fundamental decisions for 2027.

Executive agenda

Six important executive-level decisions for 2027

Each decision is grounded in either recurring evidence across our sources or a condition that materially shaped incident outcomes

The decision	Why it matters now	How you will know it worked
01 Fund phishing-resistant identity <i>For people and for machines</i>	Identity was the common denominator across the year's evidence. Misused credentials drove the largest share of confirmed incidents, appeared in every forensic investigation and were reproduced under test. Executives should treat user accounts, service principals and automation tokens as the same class of business risk: forms of access that must be strongly verified, continuously governed and quickly revoked.	→ Passkey and FIDO2 coverage is in place for privileged users, remote access and high-risk workflows. → Service principals, automation tokens, and other machine credentials are inventoried, owned and assigned expiry dates. → The organisation can show median time to revoke stolen sessions, tokens and secrets across the technology environment.
02 Open a legacy-retirement lane	Unsupported technology turns known risk into persistent exposure. It produced some of the highest-severity findings and complicated live response when incidents occurred. A patching programme cannot solve systems that no longer receive patches, so executives need a funded route to retire, isolate or formally accept legacy risk with an expiry date.	→ Unsupported critical assets are being retired, replaced or isolated, not repeatedly deferred. → Overdue critical patches are reducing month by month. → Approved exceptions have named owners, compensating controls and hard expiry dates that are reviewed by leadership.
03 Treat recovery and management systems as tier zero	The systems needed to manage and restore the business are also the systems attackers seek to control. Backup, storage, virtualisation, directory and cloud control planes either became high-value targets or critical recovery dependencies. If these assets are reachable from ordinary user environments, a single compromise can become a resilience failure.	→ Administrative consoles, backup platforms, storage, virtualisation, directory services and cloud control planes are segmented from ordinary user networks. → Privileged access is managed and monitored, while recovery tests assume identity compromise and still prove that backups can be protected, accessed and restored safely.

Continues on next page

The decision	Why it matters now	How you will know it worked
04 Govern software supply chain and pipeline trust	Trusted automation can execute attacker-controlled code with the same authority as legitimate work. This year's evidence showed how build pipelines, upstream components and runtime secrets can turn normal development workflows into attack paths. Executives should govern software supply chain trust as a resilience issue, not only as an engineering concern.	→ Build environments use pinned and verified dependencies, signed artefacts and provenance checks and workload credentials are short-lived rather than standing secrets. → Pipeline egress is controlled, and non-human identities used by CI/CD, agents and automation are inventoried, reviewed and expired.
05 Authorise containment outside business hours	More than four in ten confirmed incidents began at night or during the weekend, while regulatory and attacker clocks continue to run. Fast detection only improves outcomes if someone is already authorised to act. The board-level question is whether containment authority, escalation paths and legal thresholds are clear before the incident starts.	→ Standing containment authority is documented, approved and exercised for defined actions such as account disablement, token revocation, host isolation, blocking indicators and suspending risky access. → Detection-to-containment time is reported separately for business hours and off hours, so the board can see whether authority matches detection speed.
06 Make forensic readiness measurable	Every investigation encountered evidence gaps that limited confidence in root cause, scope or residual risk. That creates business uncertainty at the worst possible moment: when leaders must brief regulators, insurers, customers and executives. Forensic readiness should therefore be measured as a recovery capability, not treated as optional security telemetry.	→ Critical log sources are covered and retained across identity, perimeter, endpoint, cloud and workload environments. → Incident records contain stable identifiers and complete detection, escalation and containment timestamps. → In exercises and real cases, the organisation can evidence root cause, scope and residual risk rather than relying on assumption.

➔ What this means for executives

Companies should fund the work that reduces the conditions that allow incidents to escalate. First, remove attack paths where possible. Where they cannot be removed, segment them. Require strong verification for high-risk actions, and make sure anything that cannot be eliminated is properly instrumented. New tooling should only be added when it supports those decisions, not when it becomes a substitute for making them.

Executive scorecard

How to measure progress

An executive scorecard should measure reduced exposure and improved response, not activity. Seven measures are enough, and each one has an owner.

Measure	Cadence	What it tells the executives
Identity resilience	Quarterly <i>Validated by independent testing</i>	Whether a stolen credential still works. Evidenced by Phishing-resistant coverage of privileged, remote and workload identity, and median token revocation time.
Technical debt	Monthly	Whether legacy risk is shrinking. Evidenced by unsupported critical assets, overdue high-risk patches, and the age of approved exceptions.
Forensic readiness	Quarterly <i>Tested by scenario</i>	Whether you could answer a regulator. Evidenced by critical log-source coverage, retention, and timestamp completeness.
Supply chain & agent integrity	Quarterly	Whether your automation can be trusted. Evidenced by pinned dependencies, signed artefacts, and an inventory of non-human identities with expiry.
Attack path reduction	Per engagement <i>With independent retest</i>	Whether test findings translate into removed risk. Evidenced by complete paths eliminated and retest pass rate.
Response margin	Monthly <i>And rolling twelve months</i>	Whether you can act faster than an attacker moves. Evidenced by elapsed time from detection to containment, split by business hours and off hours.
Recovery resilience	Semi-annually <i>Tested by exercise</i>	Whether you could come back. Evidenced by tier zero segmentation, backup access controls, and restoration validated under compromised-identity assumptions.

Reporting cadence

01

Monthly

Operational performance, material incidents and overdue critical actions.

02

Quarterly

Attack path retesting, identity resilience, forensic readiness and legacy risk reduction.

03

Annually

Full evidence synthesis across all sources, benchmarked again against the global picture.



New tooling should only be added when it supports decisions, not when it becomes a substitute for making them.

02

Managed Detection & Response

This section looks at what reached our security operations teams during the reporting period, how often those cases became confirmed incidents, and what the timing and themes reveal about operational resilience. The important story is not alert volume, but the work required to separate routine activity from credible threat, contain incidents before impact and preserve the capacity to act when risk emerges.

In this section, confirmed incidents are counted at case-record level, as explained in the methodology. A confirmed incident is defined as a case closed as a security incident, impersonation attempt, end user data leak or customer data leak.

What reached the SOC, and what mattered

A case reaching the SOC is not the same as a threat reaching the business. Most cases were resolved through investigation, correlation or automated handling. Roughly one case in eleven required a defensive action, and one in thirty-seven was confirmed as an incident. Among the non-routine cases, seven in ten were stopped before operational impact.

The executive conclusion is that case volume alone is a poor measure of security performance. While a case record represents something requiring investigation, not necessarily a threat or security incident, most were resolved through triage, correlation, automation or investigation without evidence of operational impact.

About one case in twenty was closed by automation without analyst involvement and the other nineteen

in twenty were worked by a person. The work that matters sits between alert generation and confirmed impact: separating routine noise from credible threat, acting early enough to prevent escalation, and preserving enough analyst capacity for the cases that require judgement. A smaller incident count is therefore not evidence of low demand and may in fact be evidence that the security operation is doing its job before incidents become visible to the business.

Cases reaching the SOC are investigated before they become incidents

Cases handled by the SOC




Figure 1. Disposition of the case records handled during the evidence window.

A case reaching the SOC indicates activity requiring investigation – not that an attacker has gained access or caused impact.

➔ What this means for executives

Measure your security operation on analyst-managed case load, confirmed incidents, containment outcomes and elapsed response times. Alert volume measures tooling activity, not defensive performance. Executives should ask how many cases required human judgement, how many were contained before impact, how quickly credible threats moved from detection to action, and whether automation is reducing analyst burden without hiding risk. The aim is not to maximise alerts, but to preserve decision capacity for the cases that can affect the business.

There was no quiet season

While confirmed threats stayed inside a narrow band for the whole year, there was no month in which capacity could safely have been reduced.

The busiest full month of confirmed incidents ran to 151 and the quietest to 76. The quietest month of the year was February, not December, as might be expected. This means that holiday periods produced no meaningful relief. The more important conclusion here is that incident demand behaved like a continuous operating requirement, not a seasonal workload.

A lower month was still active enough to require sustained triage, response capacity and escalation readiness. For executives, the question is therefore not whether the SOC can absorb occasional peaks, but whether it can maintain consistent decision quality when volume does not meaningfully disappear.

Busiest and quietest months




Figure 2. Busiest and quietest months compared.

Confirmed incidents by month




Figure 2.1. Confirmed incidents by month. June 2026 is a partial month and is excluded from the median.

➔ What this means for executives

Plan detection and response capacity against a persistent annual load, not a seasonal one. The evidence does not support reducing coverage during holidays or expected quiet periods. Executives should test whether staffing, escalation paths, provider coverage and containment authority hold during leave periods, weekends and budget-constrained months. The aim is to maintain reliable response quality when there is no quiet season to recover capacity.

How attacks really arrive

Where threats could be classified reliably, more than four in ten began with a phishing message or malicious link.

Confirmed incidents by type




Figure 3. Confirmed incidents by themes out of those that could be classified.

The pattern is important because it challenges how many organisations still allocate security effort. The largest classified theme was not advanced exploitation or bespoke malware, but phishing and malicious links: routes that target users, credentials and trusted workflows. Malware and misused identity then sit close behind, reinforcing the same conclusion. The most common arrivals were not exotic; they were familiar routes that succeeded because trust could be abused quickly.

Direct exploitation was the smallest classified theme in our monitoring data, but that should not be read as evidence that exploitation is unimportant. Monitoring often sees the activity that follows initial access: credential use, malware execution, lateral movement or suspicious behaviour inside the environment. Incident response research, by contrast, can reconstruct the intrusion from its origin. Chapter 05 reads these perspectives together, because each sees a different point in the attack lifecycle.

➔ What this means for executives

Balance prevention spend against how attacks actually arrive. Do not treat exploitation as the only front door. Phishing, malicious links, malware and identity misuse point to the same priority: make trusted access harder to abuse and faster to revoke.

Detection is fast. The clock is faster.

What constrained incident outcomes was not how quickly threats were found. It was whether anyone was authorised to act at the moment they were found.

The timing is important because it changes where the risk sits. Forty-nine per cent of confirmed incidents were detected and reported in less than 20 minutes, and 30% in less than 15 minutes. Almost all were also mitigated, or at least responded to, within the same timeframe. Against a 24-hour regulatory notification window, detection and initial response are not the long part of the story. The remaining risk sits in decision-making, escalation, authority and evidence quality.

That matters because the regulatory clock starts before the organisation has finished understanding the incident. Under NIS2, essential and important entities must submit an early warning to their national CSIRT within 24 hours of becoming aware of a

significant incident, followed by a fuller notification within 72 hours and a final report within one month. Financial entities carry comparable obligations under DORA, and exact duties depend on national transposition and legal confirmation.

The finding that 43% of confirmed incidents began outside weekday business hours should therefore be read as a governance finding, not an operational inconvenience. An incident that begins at two in the morning on a Saturday does not wait for office hours before consuming regulatory time. If containment authority or escalation depends on a decision maker becoming available, the organisation can still lose time that fast detection and initial response have already won.

Timing and detection speed

When confirmed incidents began

Weekday
business hours

57%

Night
and weekends

43%

Detection speed against the 24h regulatory clock




Figure 4. Timing of confirmed incidents, and detection speed set against the NIS2 early warning window.

➔ What this means for executives

Pre-approve the containment actions that must not wait for office hours. This should include, at minimum, account disablement, token revocation, host isolation, blocking known indicators and suspending high-risk access. The decision is not whether the provider may act without oversight; it is which actions the organisation has already agreed are safe, proportionate and time-critical when the incident clock has started.

03

Digital forensics and incident response

Forensic investigation shows what happened when prevention and monitoring were not enough. Unlike monitoring, which makes clear what was detected and contained, incident response gets closest to the questions executives face after an intrusion: how did it begin, what did it touch, and what can we say with confidence? This section shows why trusted access and incomplete evidence turn technical compromise into business uncertainty.

This section uses a different evidence base from the MDR analysis before it. The previous section is based on structured SOC case records and can therefore report case-level volumes, timings and shares. This section is based on restricted forensic engagement reports, which provide deeper detail on serious incidents but cannot be reported with client counts, engagement counts or exact denominators. For that reason, the analysis describes recurring conditions using confidence language and frequency bands rather than exact case numbers.

When intrusions succeed, evidence decides

The pattern was not a single malware family or one repeated entry technique. Across the investigations the common factor was the exposure or misuse of trusted credentials, tokens and secrets, followed by activity carried out through legitimate tools, authenticated sessions, APIs or automated workflows.

What recurred across the year's forensic investigations

Every investigation	Credential, token or secret exposure	Material gaps in the available evidence
Most investigations	Substantial credential or token rotation required	Data access or exfiltration implications
Several investigations	Third-party or supply-chain trust involved	Destructive activity or persistent malware

Figure 5. Recurring conditions identified across restricted forensic investigation reports.

The exploitation of trusted credentials and tools mentioned above matters because malicious activity using trusted access is difficult to distinguish from normal operations – and much harder to scope when telemetry is incomplete.

The second finding is more uncomfortable: every investigation encountered material evidence gaps,

including missing perimeter logs, interrupted endpoint telemetry or incomplete runtime visibility in cloud and pipeline environments. When evidence is thin, organisations cannot establish root cause, bound the scope of an incident or make recovery decisions with confidence. They are forced to work from assumption rather than fact.

➔ What this means for executives

Treat forensic evidence as recovery infrastructure. Identity evidence, perimeter logs, cloud audit data and workload runtime telemetry should be owned, retained, and tested as deliberately as backup and restoration capability. Without them, root cause, scope, and residual risk remain uncertain. Executives must decide whether the organisation can answer regulators, insurers, and customers with evidence when it matters, rather than whether to buy more visibility.

Different incidents, same trust problem

While the incidents examined in our forensic engagement reports differed completely in technique, each exposed a weakness in how trust was created, retained, observed, or revoked.

How trust was established during attacks

Pattern	What happened	The executive lesson
Ransomware & domain compromise	Domain controllers were compromised, destructive activity followed, and backup infrastructure was targeted. It could not be established with confidence how the perimeter was first breached and what left the environment.	Retire unsupported edge systems, monitor privileged activity, and isolate recovery infrastructure from the domain it is supposed to protect.
Credential misuse & API exposure	A credential belonging to a former client relationship – together with cached authorisation that was never cleared – allowed continued cross-tenant access to customer information.	Ensure that offboarding revokes tokens, caches, sessions, and tenant-scoped authorisation. Remember that removing the password is not equivalent to removing the access.
Endpoint compromise & blinded telemetry	User-assisted execution triggered scripted malware, exposed credentials and created scheduled persistence. Local controls blocked repeated payload attempts, but central telemetry was partially interrupted at the same time, limiting confidence in the full sequence of activity.	Measure local prevention and central visibility as two separate things. Remove standing local administrator rights and harden script execution.
Software supply chain compromise	A trusted upstream component executed inside the build pipeline, exposed runtime secrets and transmitted encrypted data outbound. Investigators found no confirmed evidence of destructive activity or wider lateral movement.	Build pipelines should be governed as privileged environments. Dependencies must be verified, workload identities tightly scoped, and outbound traffic controlled.

Cross-case conclusion

The wider lesson is that security controls can work and recovery can remain uncertain. In several cases, controls blocked or limited parts of the attack, but the organisation lacked enough evidence to say exactly how the intrusion began, how far it reached or what risk remained. That uncertainty affects service restoration, regulatory reporting, insurance discussions and customer confidence. The issue is not a single missing tool. It is a measurable readiness gap: evidence must be retained, trusted access must be revocable, and critical environments must be observable before an incident begins.



The real readiness gap is not always whether controls work, but whether the organisation can prove what happened afterwards.

04

Offensive security

Offensive security testing means controlled security testing designed to show what an attacker could potentially achieve in an organisation's environment, not just what attackers happened to do. Its value is not in the number of findings, but in the attack paths it reveals. This section shows how familiar weaknesses combine into credible routes to impact – and why removing whole paths matters more than closing isolated issues.

This section distinguishes between scored security findings and unscored observations. Scored findings are documented weaknesses that testers assigned a severity level, such as Critical, High, Medium or Low. Unscored observations were also documented during testing, but were not assigned a severity level and are therefore shown separately. The percentages in Figure 6 are calculated only from scored findings, not from all observations recorded during testing.

What controlled testing proved

Among the findings that received a severity rating, more than one in four was Critical or High, but the main lesson is not the severity count.

The highest-risk paths combined familiar weaknesses: weak or exposed credentials, unsupported technology, insufficiently hardened directory services and management interfaces reachable from places they should not have been. Risk accumulated across the chain, not from any single finding in isolation.

The offensive evidence comes from a subset of the year's engagements and is reported as documented, without extrapolation across the wider programme. Its value is in showing how ordinary control gaps combine into credible routes to impact.

More than one in four rated findings was Critical or High




Figure 6. Severity distribution across rated offensive security findings.

Unrated exposure observations are excluded from the shares. Informational here refers to security weaknesses that are not currently exploitable and do not present a meaningful risk by themselves.

➔ What this means for executives

Measure remediation by the attack paths removed, not by the number of findings closed. A finding is only resolved when the route to impact has been broken and retested, so executives should ask which complete paths have been eliminated, which still depend on compensating controls, and which remain open because ownership or funding is unclear. Closing individual issues is useful, but it does not reduce business risk if the wider path still works.

Four failures that keep returning

The recurring weaknesses found in offensive security testing were structural rather than exotic: weak and exposed credentials, unsupported systems, weaknesses in directory services and overly reachable management planes. All four sit inside a small set of enterprise control domains that an executive owner can be held accountable for.



Identity

Weak and exposed credentials accelerated access

Password spraying, credential reuse, weak service accounts and stolen sessions created working routes to initial access and to escalation.



Directory services

Identity configuration weakened containment

Incomplete local administrator password management, weak signing controls, Kerberoasting exposure and privileged account gaps all increased the reach of a single compromise.



Technology lifestyle

Unsupported systems remained exploitable

Known vulnerabilities and obsolete platforms produced remote code execution and privilege escalation opportunities that no amount of monitoring would have prevented.



Management planes

Reachability converted footholds into impact

Administrative consoles and infrastructure platforms were reachable from broad internal segments, and in some cases from the internet.


➔ What this means for executives

Assign named ownership for identity, technology lifecycle, directory services and management planes. These recurring failures are not isolated technical issues; they are control domains that shape whether a single weakness becomes a route to impact. Executives should track whether exposure is being removed, management access is being segmented, privileged access is being reduced and fixes are passing retest.

The treatment model

There are four ways to treat an attack path, and they should be applied in order. Remove the condition where possible. Where it cannot be removed, segment it. Where access must remain, verify high-

risk actions. Detect what is left. Many organisations start at detection because it is the easiest response to buy, but visibility is not the same as risk reduction. If the path still exists, the business is still exposed.



Step	What good looks like	Evidence of progress
Remove	Unsupported assets, dormant accounts and obsolete services are decommissioned rather than documented.	Assets and identities eliminated, counted and verified.
Segment	Management interfaces and administrative protocols are restricted by tier.	Verified blocked paths from ordinary user networks.
Verify	High risk actions require phishing-resistant authentication and stronger process checks.	Coverage figures and the results of bypass testing.
Detect	The paths that cannot be removed or isolated are fully instrumented.	Telemetry coverage and detections proven by test.

➔ What this means for executives

Ask every security investment which treatment step it supports: remove, segment, verify or detect. If the answer is repeatedly detect, the organisation may be buying visibility into risks it could have reduced. Detection matters, but it should not substitute for removing obsolete assets, segmenting management access or reducing privileged trust. The question for executives is whether the attack path is being reduced, not whether it is better observed.



05

The recurring threats across our evidence

The strongest findings in this report are the ones
that appeared from more than one angle.

Where the evidence points in the same direction

The previous sections looked at each evidence source separately. Managed Detection & Response showed what reached our SOC and became confirmed incidents, digital forensics and incident response showed what investigators found after serious intrusions, and offensive security testing showed which weaknesses could be reproduced under controlled conditions.

This section in turn brings those three measurements together. Its purpose is to identify the threats and control weaknesses that appeared across monitoring, investigation and testing – not just once, and not from

one perspective alone. When the same condition is seen in live operations, post-incident investigation and controlled testing, it becomes one of the report's strongest signals: a recurring threat that is real, repeatable and important enough to shape executive-level decisions.

The table below therefore focuses on the conditions that appeared across all three evidence sources. These are the threats and weaknesses that should matter most to leaders because they are not isolated findings; they are the patterns most consistently visible across the year's evidence.

The condition	What monitoring saw	What investigation found	What testing proved
Trusted access can be abused	Phishing and misused valid accounts were the largest classified themes.	Every investigation involved credential, token or secret exposure.	Weak passwords, session compromise and directory gaps enabled escalation.
Old technology keeps known risk alive	Exploitation-related cases were present but smaller than identity themes.	Unsupported perimeter infrastructure was an assessed access path in a destructive incident.	Obsolete platforms and exposed services produced Critical and High findings.
One foothold can reach critical systems	Remote access and lateral movement behaviour required investigation.	Domain controllers, APIs, cloud applications and pipelines widened the trust boundary.	Management interfaces were reachable through excessive internal trust.
Weak evidence turns incidents into uncertainty	The case export lacks stable identifiers and containment milestones.	Every investigation met material telemetry or evidentiary limitations.	Coverage held where telemetry existed and failed where it did not.

Two important conditions appeared in only one evidence source, and that is part of the lesson. Incident timing appears mainly in monitoring because testing happens in agreed windows, while attackers choose their own timing. Supply chain compromise appears mainly in investigation because it often sits

outside the scope of conventional testing. Neither condition is less important because it appeared once. Instead, they show why the report uses three evidence sources: each sees a different part of the risk, and together they give leaders a fuller basis for decision-making.

➔ What this means for executives

The greatest exposure sits where trusted access, unsupported technology and broad internal reachability meet. Any one of these may look manageable on its own. Together, they create attack paths that can spread quickly. Where evidence is thin, the organisation then faces a second problem: it cannot confidently explain what happened, how far it reached or what risk remains.

06

The wider picture

No single provider sees the whole threat landscape. We therefore read our European operational evidence alongside major public threat research from the same period to understand which findings are specific to the environments we monitored, which are part of a wider industry pattern, and which should shape executive-level planning for the year ahead.

What the industry saw, and what we saw

The clearest overlap between our evidence and wider industry research is identity. Independent datasets, built on different populations and methods, point to the same conclusion as our own evidence: credentials are now central to modern intrusion activity.

Public breach research places credential abuse in roughly two of every five breaches when the full attack chain is counted rather than only the first step, while industry telemetry shows that most detections now involve no malware at all. Attackers log in, then use the tools already available to them. Our monitoring, investigations and testing each found that pattern separately, which is why identity leads the executive agenda in this report.

A second pattern is forensic uncertainty. Industry research shows that long-running intrusions can

persist for months, often beyond the log retention many organisations keep. When that happens, investigators may be unable to prove how the intrusion began, what was touched or what risk remains.

Our own investigations encountered the same condition. This makes forensic readiness more than a technical improvement; it becomes an executive funding priority because leaders need evidence they can use with regulators, insurers, customers and internal decision-makers.

Industry findings	Our monitoring	Our investigation	Our testing
Credential and identity abuse runs through the whole attack chain	Confirmed	Confirmed	Confirmed
Attacks increasingly succeed without malware	Confirmed	Partly confirmed	Confirmed
Evidence and log-retention gaps limit investigation	Confirmed	Confirmed	Partly confirmed
Exploitation of internet-facing systems leads initial access	Partly confirmed	Confirmed	Confirmed
Ransomware targets backup, identity and virtualisation	Not observed	Confirmed	Partly confirmed
Third-party and supply-chain compromise is rising	Not observed	Confirmed	Not observed
Social engineering is moving off email to voice and chat	Not observed	Partly confirmed	Not observed
Attackers are using AI to accelerate existing tradecraft	Partly confirmed	Partly confirmed	Not observed

Figure 8. Prominent findings from 2026 global threat research, read alongside our own three evidence sources.

➔ What this means for executives

When a finding appears both in our own operational data and in independent global research built on different methods, it stops being a vendor observation and becomes a planning assumption. That is the threshold leaders should use when deciding what to fund, what to measure and what to challenge.

Three insights from the industry, read against our year

The previous page showed where our evidence and global research pointed in the same direction. This one looks at what the wider industry helps us see differently.

Risks that monitoring may see later, controls that sit outside traditional telemetry, and attacker timelines compressing the space between detection and action. For executives, the lesson is not to choose one evidence source over another, but to understand what each source can and cannot reveal.

Insight 01

The edge is where intrusions begin

Multiple independent 2026 datasets now place exploitation of internet-facing systems as the leading initial access vector, at roughly a third of intrusions. In our monitoring data, exploitation was the smallest classified theme, but that does not make the two observations contradictory.

They describe different points in the same attack lifecycle. Incident response research often reconstructs an intrusion from its origin. Monitoring watches continuously and usually sees the intrusion later, once activity has moved inside the environment.

Our other evidence sources sit closer to the industry view. Unsupported perimeter infrastructure was an assessed access path in a destructive incident, and obsolete platforms produced Critical and High findings under test.

Read together, the lesson is clear:

Edge infrastructure should not sit outside security governance simply because it is outside the endpoint

technology environment. It must be included in monitored scope, patching governance, vulnerability management and incident response planning.

01 The practical lesson

It's about scope rather than tooling. Detection coverage is only as good as the surface it can observe, and internet-facing infrastructure is often outside the technology environment a monitoring service has been given access to.

Leaders should confirm, deliberately, that edge devices are covered by monitoring, patching governance, vulnerability management and incident response playbooks.

Continues on next page

Insight 02

Social engineering has moved off email

Voice phishing rose sharply in global incident response data during the period and is now among the most common initial vectors, particularly for cloud compromises.

Pretexting has also been formalised as a distinct vector in breach research, and a significant share of social engineering breaches now arrives through channels that email gateways never see.

A phone call to a service desk leaves no artefact in security telemetry. That makes this a governance problem, not only a detection problem.

02 The practical lesson

Organisations need service-desk identity-verification procedures, callback policies, clear staff authority to refuse unsafe requests, and testing that reflects how social engineering now works.

Insight 03

Attackers are trading access at speed

Industry research indicates that the hand-off from an initial access broker to the group that causes the damage can now happen in seconds rather than hours. That changes the meaning of a low-priority alert. A commodity infection that once might have sat in a queue can become the beachhead for a targeted operation almost immediately.

03 The practical lesson

The lesson for executives is that early-stage detections should not be treated as low-consequence simply because impact has not yet occurred.

Access now has a resale value, and the time between initial compromise and operational risk is shrinking. Detection therefore needs to be paired with pre-agreed containment authority.

➔ What this means for executives

The wider industry picture does not replace our own evidence, but explains its blind spots. While monitoring shows what becomes visible inside the environment, the industry data makes clear why leaders must also govern the edge, test social engineering beyond email, and reduce the time between detection and containment. The priority for 2027 is not more visibility in isolation. It is ensuring that the right surfaces are in scope, the right decisions are pre-authorised and early signs of access are treated before they become impact.



07 Outlook

No report can predict the whole threat landscape for the year ahead. We therefore use the patterns identified across our European evidence and wider industry research to consider which risks are likely to become more important in 2027. This section focuses on where familiar weaknesses may become more consequential as AI, automation and machine identity expand.

AI did not change the attack. It changed the pace.

The most important finding about AI this year is a reassuring one: rather than producing a new category of attack, it has made the existing ones faster, cheaper, and more numerous.

That conclusion is consistent across the year's research and matches what we saw in practice. The incidents we investigated and the intrusions we detected still relied on familiar techniques.

That is exactly what the industry data suggests:

AI is not replacing established attacker tradecraft but making it easier to scale. For decision makers, the distinction matters. This is not a call for a new security strategy, but for faster execution of the one organisations already need.

What the evidence shows

The same attacks, delivered at greater scale

Where researchers have analysed adversary use of AI directly, the assisted activity maps onto familiar techniques: mostly phishing, then vulnerability exploitation. The craft is the same. The rate has changed.

What users face

The old phishing warning signs are disappearing

European research estimates the large majority of phishing campaigns now use AI-generated or AI-enhanced content. Awareness training built around poor grammar and clumsy formatting is training for a threat that has retired.

Where the pressure lands

The response window is getting smaller

Industry telemetry attributes part of the collapse in attacker timing to AI assistance across reconnaissance, credential theft and evasion. Response tempo, not control design, is where that pressure lands.

What comes next

Malware is starting to adapt as it runs

Researchers have now documented malware families that call large language models while running, in order to regenerate their own behaviour and evade static detection.

➔ What this means for executives

AI does not make existing controls obsolete and the techniques documented so far are still addressed by the same fundamentals: phishing-resistant authentication, removed attack paths, segmented management planes and fast containment. What changes is timing. Because AI makes familiar attacks easier to scale and harder to spot early, the six decisions on page 8 of this report become more urgent, not less relevant. The answer is therefore not a separate AI security strategy, but faster delivery of the controls organisations already need.

AI agents will turn trust into the next attack surface

The previous section addressed attackers' use of AI. This section addresses a different risk: organisations deploying AI agents as trusted actors inside their own environments.

The biggest AI-related exposure for 2027 may not be attackers using AI. It may be organisations deploying AI agents faster than they can govern them. This year's evidence already showed the shape of that risk: trusted automation with standing access, limited oversight and weak revocation. AI agents will make that problem more common and harder to contain unless organisations act before adoption accelerates.

One investigation this year from our external sources involved a trusted automated pipeline that executed an upstream component and exposed the runtime secrets it held. It was not an AI incident, but it showed the same pattern organisations should expect from poorly governed AI agents: a non-human identity, trusted by design, able to reach real systems, holding credentials and acting without direct human supervision.

The governance failures seen in that case are the same failures that could make AI agent deployments risky at scale: credentials that do not expire, permissions granted once and never reviewed, limited runtime telemetry and no fast way to revoke access. The lesson is not that AI agents create an entirely new problem, but that they can amplify existing weaknesses in non-human identity, access governance and trusted automation. Organisations that fix those weaknesses now will be better placed to adopt AI agents without creating a new route to incident impact.

What we expect to matter, and why we think so

- Non-human identities will grow faster than identity governance. Agent adoption will increase the number of machine actors with access to enterprise systems, but many organisations still govern identity as if employees are the only users that matter.
- Over-permissioning will matter more than exotic model attacks. Poorly scoped access is an old failure in a new place, which means the controls already exist if organisations choose to apply them to agents.
- Prompt injection will be managed through containment, not solved outright. Because instructions and data share one channel, organisations should limit what agents can do, require approval for high-risk actions and retain the ability to stop them quickly.
- Unsanctioned agents will appear before formal governance does. They are likely to arrive through developers and business teams before procurement or architecture review, which means leaders should assume some agent use is already outside the official inventory.
- AI platforms themselves will become targets. Attackers are already treating deployed AI tooling as part of the attack surface, including the systems agents rely on and the instructions they can be made to follow.

How to prepare for 2027

Before expanding AI agent use, organisations should inventory their non-human identities, replace standing credentials with short-lived and scoped access, and ensure every agent has a working kill switch. Agent activity should be logged to the same standard as privileged employee activity. Leaders should then rehearse the scenarios this year's evidence points towards: supply chain compromise, mass secret rotation, session revocation, backup targeting and loss of endpoint telemetry.

The risks already taking shape for 2027

Each of the below items is anchored in either our own evidence or independent research from the same period.

Machine identity

Attackers will follow the secrets

Pipelines, cloud applications, service principals and automation tokens can turn one compromise into wider access far faster than the theft of a single user account. Yet they are rarely governed with the same discipline as employee identities, as was seen in our own data.

Software supply chain

Trusted pipelines will become the breach path

Malicious packages in public repositories have grown by orders of magnitude in recent years. When dependencies can change after they are trusted, an upstream compromise can move through many environments before anyone has reviewed it.

Recovery as a target

Attackers will target the way back

Global research shows ransomware operators increasingly trying to destroy the route back to normal operations by targeting backups, identity services and virtualisation together. One of our own investigations followed the same pattern.

Session and token abuse

The login may end, but the access remains

Cached tokens, long-lived sessions and standing grants can keep access alive after a password is changed or a user is removed. During the year, we saw how revocation can fail when access is spread across more than one trust store.

Leading indicators worth putting on a dashboard

01

Growth in long-lived machine credentials and secrets available to automated workloads and AI agents.

02

Use of mutable tags, unverified dependencies, and publicly reachable pipeline or cluster control planes.

03

Percentage of critical systems without complete identity, perimeter, cloud and runtime telemetry.

04

Time required to revoke sessions, tokens and workload credentials across every trust store you own

05

Number of tier zero systems reachable from user or general server networks.

06

Elapsed time from detection to containment, measured separately for business hours and off hours.



2027

The biggest AI-related exposure for 2027 may not be attackers using AI. It may be organisations deploying AI agents faster than they can govern them.

08 Appendix

Method, scope and limitations

We publish our methodology to make the basis for the report's findings transparent and open to scrutiny.

Evidence window

Monitoring data covers the period from 1 July 2025 to 11 June 2026. June 2026 is a partial month, is identified as such in all relevant charts, and is excluded from averages and medians. Forensic and offensive security evidence is drawn from the same reporting period.

During this evidence window, Conscia MDR protected 71 organisations. The monitoring evidence was generated through Conscia's European SOC in Slovenia and supported by our SOC capabilities across Denmark, Spain and the Netherlands.

In this report, European data means evidence generated from the European environments monitored, investigated or tested by Conscia during the evidence window. It does not mean a statistically representative sample of every European market, sector or organisation.

We therefore use the data to identify patterns observed in our operational work, and compare those patterns with wider research before drawing board-level conclusions.

Managed Detection & Response evidence

→ Case records are drawn from our security orchestration platform across the European environments monitored by Conscia's SOC and MDR teams during the reporting period. Across the evidence window, 260,943 investigations

were performed. Of these, 1,266 were confirmed as security incidents and successfully contained.

- A confirmed incident is defined as a case where active attack bypassed automated security-tool response logic and was prevented or contained by Conscia MDR. No Conscia customer experienced harm from a Category 5 cyber incident during the evidence window.
- The export does not include a stable cross-case incident identifier; totals are therefore reported at case level. Where two cases arise from one intrusion, they are counted separately.
- Time to detect is the period from when an alert is generated until the case is classified as either benign or malicious.
- Off hours refers to weekends and to weekdays before 08:00 or at and after 16:00, Central European Time.

Forensic and offensive evidence

- Investigations are drawn from restricted engagement reports, some of which carry handling restrictions that permit only anonymised, aggregated conclusions. Statements distinguish between confirmed, assessed, possible and undetermined findings.
- Evidentiary depth varies across the investigation set. Some findings are based on extensive forensic analysis, one is based on a high-level endpoint review, and one relies predominantly on client-provided information without independent verification. Conclusions are weighted accordingly. Frequency bands describe only the investigations supplied for this edition.
- Offensive findings are drawn from a subset of the year's engagements and are counted as documented, without reweighting. Unrated observations are excluded from all severity shares, and no engagement-level success rate is inferred.

External comparison

- Chapter 05 compares our European operational evidence with major public threat research covering the same period, including reporting from the European Union Agency for Cybersecurity, the Verizon Data Breach Investigations Report, Mandiant M-Trends and CrowdStrike Global Threat Report analysis.
- These datasets use different populations and counting rules from ours and from each other. They are used to corroborate direction and identify blind spots, not to calculate a combined figure.

What we withhold, and confidentiality

- Client identities, client counts, case counts and engagement counts below regional level are not reported. Where a denominator could reveal sensitive information, only the proportion is provided.
- Names, accounts, credentials, domains and hostnames are excluded throughout, and all source material remains restricted.
- This edition establishes the 2026 baseline. Trend analysis will begin in 2027, once stable definitions and containment timestamps are incorporated into the data model.



About the report's author

David Kasabji

David Kasabji is the Head of Threat Intelligence at Conscia. He leads the development and delivery of actionable intelligence across cyber defense and managed security operations, translating complex threat activity into clear outcomes for different audiences – from SOC analysts and incident responders to executive stakeholders and external communications.

His work spans end-to-end intelligence operations: collection and analysis of adversary activity, threat actor and campaign profiling, IOC and TTP development, and intelligence-driven guidance for detection, threat hunting, and security prioritisation. David is also actively involved in Digital Forensics and Incident Response, supporting investigations and crisis situations with rapid triage, context, and strategic recommendations. A strong focus of his role is continuously improving how intelligence is operationalised through standardisation and automation to ensure it is timely, relevant, and measurable.

Glossary

MDR (Managed Detection and Response): A cybersecurity service that continuously monitors an organisation's systems for threats, investigates suspicious activity and helps respond to incidents.

SOC (Security Operations Centre): A team and set of technologies responsible for monitoring, investigating and responding to cybersecurity threats and incidents.

FIDO2: An open standard for secure, passwordless authentication using cryptographic credentials rather than passwords.

Passkey: A passwordless sign-in credential stored on a device or password manager that uses public-key cryptography and typically a fingerprint, face scan or device PIN to authenticate the user.

DORA (Digital Operational Resilience Act): An EU regulation that sets requirements for financial organisations to manage, test and strengthen their digital operational resilience and cybersecurity.

Edge systems: Devices, systems or services located at the edge of an organisation's network, often connecting directly to users, devices, the internet or external networks.

Kerberoasting: An attack technique in which an attacker requests service-account authentication tickets from Active Directory and attempts to crack them offline to recover account passwords.

Trust Store: A system or location that stores information used to decide which identities, tokens, sessions or certificates should be trusted.



Conscia
Secure progress

Cyber Resilience Insights 2026