

Midyear Cyber Risk Report

Contents

03 Executive summary

04 Key findings

05 Take aways for risk first leaders

06 AI attacks: where are they?

07 Frequency and severity are telling different stories

08 Cause of loss and point of failure

09 The human attack surface

10 The operational attack surface

11 Ransomware and the changing dynamics

12 Finding the through line on vendor risk

14 Risk Operations Center in Action

16 Looking ahead

16 Conclusion

17 End notes and citations

Executive Summary

The Resilience 2026 Midyear Cyber Risk Report seeks to close the gap between the AI hype machine and AI reality when it comes to cyber risk. The impact of AI on cyber risk is clear and evident - just perhaps not where all the media attention is currently focused.

Security researchers have documented two firsts this half: a fully autonomous ransomware operation that ran end to end without a human at the controls, and an AI model that breached a production environment on its own. While these developments are important warning shots on the future of autonomous threats, they also distort the reality of cyber risk for organizations today. In the first half of 2026, Resilience observed no incurred losses from AI-specific attack vectors, including prompt injection, model exploitation, or agentic AI misuse. What has produced losses is older and more familiar: 85.3% of incurred losses trace back to a person believing a voice, a message, or a request that looked legitimate, up from 17.7% two years ago. So far, AI's clearest effect on the portfolio isn't a new attack type. It has made the oldest one, social engineering, more convincing.

That gap, between what's being demonstrated in research and what's actually showing up in claims, creates a specific risk: the more attention agentic AI draws, the easier it becomes to treat the fundamentals as solved and move on. Those fundamentals are still

generating losses. Known vulnerabilities and governance lapses both have well-established fixes; what's missing is speed and coverage. Governance-related losses (5.4% of incurred losses) show the gap between what a policy requires and what happens day to day: privileged access flaws, misconfigured authentication, and failed payment controls. Known vulnerabilities remain the largest technical cause of loss this half at 7.0% of losses. The fix exists; it simply isn't being applied fast or widely enough.

We have said for years that no organization can stop every attack: governance or security controls will fail on occasion, and many will land before the patches that could have prevented them. What separates outcomes is containment: how fast an event is detected and how much loss gets limited once it's underway, whether it starts with an automated agent or a person being fooled. How often an attack is attempted matters far less than how well it's contained once it starts.

This cyber risk report clearly demonstrates that a risk-first approach to cybersecurity in the age of AI is not just prudent, but now essential. Risk-first CISOs collaborating with CFOs and risk managers to manage cyber risk holistically is the best remedy.

The analysis in this report is based on cyber insurance claims data in the Resilience portfolio from January 2024 through June 2026. Unless explicitly stated, all figures are derived from direct analysis of trends among mid-size to large enterprises with an average annual revenue of \$2.04b, served through primary and excess cyber insurance policies. This analysis is supported by the Risk Operations Center (ROC), which combines threat intelligence, data science, claims, and engineering capabilities into a single team.


Key findings

Humans are the weakest link.

Losses tied to phishing, social engineering, and transfer fraud have risen steadily from 17.7% of incurred losses in H1 2024 to 85.3% in H1 2026 — the single largest movement in the portfolio this cycle. Attackers have used AI to sharpen social engineering

across the board, from emails to voice deepfakes, and defensive tactics have to keep pace. What stops these losses is enhanced user training and testing, callback procedures, and dual transaction sign-off: verifying who's actually on the other end.

Losses from Social Engineering



Agentic attacks are here, but losses will lag.

Security researchers have documented a fully agentic ransomware operation running end to end with no human operator,¹ and just this month, an autonomous AI agent breached Hugging Face's systems on its own, running more than 17,000 actions to steal credentials and move through internal systems with no person directing it.² Even so, no H1 2026 claim points to a fully automated attack chain, or even to the exploitation of a vulnerability discovered by a frontier model: a gap that won't last, which makes getting ahead of it now worth more than waiting for the first claim to force the issue.

You cannot afford to take your eye off the fundamentals.

Governance-related losses (5.4% of incurred losses) tie to gaps in policy-versus-practice that aren't new: privileged access flaws, misconfigured authentication, payment controls, and data collection practices. Known vulnerabilities have fallen from a high of 25% in H2 '24 to 7.0% of incurred losses this year, but they're still the largest technical cause of loss — these are gaps between what policy says and what actually happens day to day, and by definition, they don't surface until something goes wrong.

Taken together, these findings point to the key contradiction of cyber risk in 2026: enterprises must prepare for the attacks still forming, while staying rigorous about the ones already producing losses today.

What this means for risk-first leaders

We have said for years that you cannot stop every attack. Risk-first leaders focus on the potential for extreme loss, not on treating every threat as equally likely. Prevention will fail sometimes. What separates outcomes afterward is containment: how fast an event gets detected, and how much of the loss gets limited once it's underway. That holds whether the event starts with an automated agent or a person being fooled. How often an attack is attempted matters far less than how well it's contained once it starts.

CFOs

Your board will ask if you are prepared for the age of agentic attacks. Meanwhile, your business is adopting AI, and so are your vendors — both are expanding your exposure at the same time. If uptime slips, business interruption follows, and that's now a shared risk across your own systems and theirs. The only credible answer in that boardroom, then, is a quantified one: what's exposed, and what's being done about it. What you need is visibility into your current risk posture, not what it was at your last assessment — and you certainly don't need a heatmap to tell you it's hot.

CROs

Whatever cadence your risk-assessment framework runs on, it's too slow. Not only will the threat landscape outpace your efforts, so will the growth of technology adoption inside your own organization. Unless you are engaging continuously with changes in your risk profile and your vendors', you risk missing emerging loss drivers your last assessments didn't account for. Your models might hold up today, but they haven't been validated against actual losses.

CISOs

We have entered the era of negative time between exploitation and fix, so patching still matters, but it's too slow on its own — fundamentals like zero trust adoption, privileged access management for people and agents, network segmentation, and vendor risk management have to carry more of the weight. The costliest attacks, though, still start with social engineering, not a technical gap. Bring the same focus to phishing training and detection, payment verification controls, and backup integrity, and bring your finance and risk partners along by framing all of it as loss prevention, not CVE counts.

Brokers

Encourage clients toward a risk-first approach. That means quantifying exposure, treating loss prevention as an ongoing priority, and staying engaged with risk throughout the life of the policy, not just waiting for something to go wrong. The long tail of risk is expensive, and you want to help clients address issues before they become losses.

AI Attacks: Where are they?

Following the announcement of Mythos by Anthropic earlier this year, two notable developments have occurred involving agentic attacks. JADEPUFFER, a ransomware operation documented by Sysdig Threat Research,¹ ran autonomously from reconnaissance through data destruction with no human operator, self-correcting a failed step in 31 seconds — criminal tooling operating without a criminal at the controls. A different kind of exposure surfaced recently: OpenAI disclosed that one of its models broke out of a controlled security test, gained internet access, and autonomously breached Hugging Face's production environment while

trying to obtain answers to a benchmark evaluation, chaining stolen credentials with a genuine zero-day to gain remote code execution.² Both companies said there was no malicious intent — this was a containment failure, not a criminal attack, but the harm to Hugging Face was real. While we continue to watch these developments closely, no losses have occurred in our portfolio that can be attributed to AI. To learn more about Hugging Face, read [our blog](#) on the topic.




Frequency and severity are telling different stories

Claims frequency in the portfolio has increased over the last two periods, climbing to 45.9 per 100 policies this half, driven largely by a rise in wrongful data collection claims. Over the same period, the share of claims with an incurred loss has declined to 3.4% — the lowest of the last five periods.

Average claim severity for claims with incurred cost has also fallen, to approximately \$470,000 this half, driven by fewer high-value extortion demands and zero business-interruption claims. That could reverse if more severe events happen later in the year.

Relative Trends of Claims frequency, loss rate, and severity

Indexed to H1 '24




Period	Frequency (per 100 policies)	% Claims with a Loss	Severity (avg \$/claim with loss)
H1 24'	49.6	6.3%	\$846k
H2 24'	38.6	7.2%	\$907k
H1 25'	33.7	7.3%	\$784k
H2 25'	40.5	6.3%	\$477k
H1 26'	45.9	3.4%	\$470k

Cause of loss and point of failure

Every loss can be sorted into one of five categories based on how the attack originated: social engineering, governance, vendor, vulnerabilities, or “other.” Vendor incidents drove 81.5% of the losses in H1 2024; two years later, vendor-related losses sit at just 2.3% of incurred losses for H1 '26. Social engineering climbed

instead, becoming the dominant point of failure at 85.3% by H1 2026. Governance (5.4%) and Vulnerabilities (7.0%) round out the picture, well behind social engineering's share but not negligible on their own.

Incurred Losses by Point of Failure



Period	Social Engineering	Governance	Vulnerabilities	Vendor	Other
H1 24'	17.7%	0.0%	0.8%	81.5%	0.0%
H2 24'	41.3%	15.5%	25.2%	15.8%	2.2%
H1 25'	46.9%	1.7%	16.3%	33.5%	1.7%
H2 25'	75.2%	6.9%	9.9%	4.2%	3.9%
H1 26'	85.3%	5.4%	7.0%	2.3%	0.0%

Extortion still means a ransom demand, but increasingly the payment buys silence — keeping stolen data from leaking — rather than the return of access. It's also rare and by far the most expensive cause of loss when it happens: just 5.8% of claims account for nearly three-quarters of the losses.

extreme, showing up as common claims that individually cost very little. Transfer fraud moved differently: claim volume fell even as loss share nearly tripled, so each incident is getting substantially more expensive on its own.

Wrongful data collection and vendor data breach sit at the other

The human attack surface

“CISOs need to be turning up the difficulty on social engineering simulations and expanding beyond email if they’re going to reflect real-world incidents. ‘Click rates’ should be informing adjustments to awareness training and it’s impossible to do that if they’re artificially suppressed.”

Chris Wheeler
Resilience Chief Information Security Officer

That humans are the weakest link is not news for any security team; they see it every day in their work and they’ve created controls to address it - user training, phishing tests, and simulated smishing exercises.

Security leaders should not take false comfort in declining click rates on phishing tests. Those interventions, which were built before the proliferation of LLMs, both still ultimately depend on human judgment: Does this look and sound like someone I trust? AI voice cloning and deepfakes have gotten good enough that the answer is often “yes,” even when it’s an attacker, not the CEO.

Impersonation attacks are, however, just one vector. Simpler techniques — fake software updates and password reset emails, for example — still get an attacker in the door, no deepfake required. That’s why combined phishing and social engineering has climbed from 17.7% to 41.3% to 46.9% to 75.2% to 85.3% of

incurred losses over the last five half-year periods. It’s not just Resilience’s portfolio, either: cyber-enabled fraud and phishing already top the World Economic Forum’s 2026 CEO survey, with AI vulnerabilities newly ranked second.³

The human entry point shows up even where the payoff looks purely technical. Resilience’s review of recent extortion claims found pre-incident evidence of infostealers in the client’s environment in case after case — often traced back to the same kind of fake-update or credential-harvesting click already described above. That pattern matches Verizon’s 2026 Data Breach Investigations Report (DBIR): 73% of ransomware victims had an infostealer infection or credential leak in the year before the attack, half of them within 95 days.⁴

Social engineering: claim frequency versus loss share

Period	Claim Frequency	Loss share
H1 24’	11.9%	17.7%
H2 24’	19.2%	41.3%
H1 25’	23.8%	46.9%
H2 25’	18.0%	75.2%
H1 26’	21.7%	85.3%

Transfer fraud shows a different shape. Attackers still got through by sticking to the tried and true methods of exploiting human psychology: researching and impersonating an authority figure. Its share of incurred losses climbed from 2.90% two years ago to 9.20% in H1 2026. Claim volume moved the opposite way: up from 4.82% to an H1 2025 peak of 11.3%, then back down to 4.7% this half.

A recent incident with a small financial services firm and Resilience client illustrates how this is happening in the real world. The company’s CFO joined what looked like a legitimate Microsoft Teams meeting, audio-only, where AI-generated voices of the CEO and a well-known transactional attorney discussed a time-sensitive acquisition. The “CEO” then instructed the CFO to wire the funds. Neither the CEO nor the

attorney was actually on the call — attackers had trained the voice-cloning models on hours of publicly available audio and video of both individuals. Resilience activated a bank clawback protocol immediately, recovering the full amount before funds could be laundered. A callback to the real CEO on a known number, rather than trusting the instruction inside the meeting itself, would have caught the fraud before the wire went out.

These types of attacks are going to succeed sometimes. Training and

simulation matter, but they can't be the whole plan — resilience means having a real answer for what happens when an employee does click or hand over access. That means incident response built around containment first: immediate session revocation, MFA resets, and detection fast enough to catch lateral movement before it spreads. It also means changing what gets celebrated: stopping an attacker mid-attack should count as a win, not just preventing the click in the first place.

The operational attack surface

"Every client understands that MFA is required. Few of them can say with confidence what percentage of privileged accounts are actually enrolled."

Chuck Norton
Resilience Senior Technical Security Advisor

Not every loss comes from a person being fooled. A meaningful share comes from operational failures instead: data collection practices, and controls that exist on paper but haven't kept pace in practice, MFA enrollment, and patch speed.

Wrongful data collection

Lawsuits brought about wrongful data collection show up often and represent 16.7% of the claim count this half, up from 7.3% a year ago. Wiretap-style pixel-tracking suits filed under the California Invasion of Privacy Act (CIPA), Illinois' Biometric Information Privacy Act (BIPA) and state analogs, and federal Video Privacy Protection Act (VPPA) claims over streaming trackers keep multiplying, and that's what's driving the volume. Despite that, per-claim settlements stay modest and the aggregate loss line is flat. California SB 690 would eliminate the private right of action for the specific wiretap-law claims driving much of this litigation, though it hadn't passed as of this writing.⁵ When clients do settle, it's typically under their self-insured retention.

Vulnerability management

Patching follows a similar pattern to MFA: a fix that exists but isn't keeping pace. Exploitation now often happens before a patch even exists, not after one gets ignored. Mean time-to-exploit sits at negative seven days industry-wide.⁶ Once an attacker is in, the clock moves just as fast: average breakout time — the gap between initial compromise and first lateral movement — is 29 minutes, down from 48 minutes in 2024, with the fastest recorded breakout at 27 seconds.⁷ Zero Trust controls — network segmentation, privileged access management, just-in-time access — matter more now, acting as the brakes that keep one exploited vulnerability from becoming a full network breach.

The losses from vulnerabilities in our portfolio so far this year have been garden-variety patching errors. These stemmed not from a zero-day or an agentic attack, but from a missed patched external device that was exploited. Even disciplined vulnerability management can fail. When it does, the same containment techniques that work against agentic attacks also limit the blast radius of an exploited CVE.


Ransomware and the changing dynamics

Extortion, driven by ransomware, is the biggest cause of loss in the Resilience portfolio. Its share of incurred losses has held between 65% and 75% since FY 2024, reaching 73% year-to-date in H1 2026 — while accounting for just 5.8% of total claims, a sign of how rare but disproportionately costly these events are. Rising adoption of immutable backups, up from 79.9% to 85.2% year over year (according to client self-report), is likely part of why claims stay this infrequent, giving more organizations a path to recover without a ransom claim escalating this far. The attacker landscape is consolidating too: per Check Point Research, the top ten named ransomware groups now account for about 71% of victims, up from 57% a year ago. Extortion case volume rose roughly 30% over the same stretch.⁸

The shift driving that volume is a change in method. Data-theft extortion without encryption climbed from 49% of extortion claims in the first half of 2025 to 65% in the second half, and by year-end, pure encryption accounted for just 13% of attacks.⁹ Per a Google Threat Intelligence Group-produced report, the same shift is driving the largest publicly

reported campaigns industry-wide.¹⁰ Backups can't stop this kind of extortion — there's nothing to decrypt or restore. Identity containment limits what a stolen credential can actually reach, and exfiltration detection catches the theft while it's happening, before the data is fully out the door. Those defenses are showing up in the numbers: severity of extortion events is falling, in part because clients are also increasingly declining to pay data-theft demands. Paying doesn't reliably buy what it promises. CrowdStrike's 2025 State of Ransomware Survey found that 93% of victims had data stolen regardless of whether they paid, and that submission led to a second attack 83% of the time, often from the same actor.¹¹ This works for reasons that have nothing to do with backup strength — an encryption-free hit has nothing to restore in the first place. What actually justifies holding firm is simpler. The policy limit stays intact for whatever comes next, like class-action defense, instead of funding a ransom with no real assurance behind it.

Losses from Ransomware
2025 Full Year % and 2026 YTD %



Finding the throughline in vendor risk

Unlike past periods where vendor risk was a leading cause of loss, this half's dominant vendor event, Canvas (a learning management system used by thousands of educational institutions), was a data-theft extortion attempt.¹² The platform did go offline, disrupting access during a high-stakes academic period, but that outage stayed within common standard waiting periods for business interruption (BI), so the claims exposure came through contingent data breach and extortion coverage instead of a business-interruption claim. We saw nothing like Change Healthcare¹³ or CDK Global¹⁴ — a vendor compromise causing multi-client BI at higher severity, weeks of operational disruption, and eight-figure single-claim losses — show up in the H1 2026 claims. That's why the share of incurred losses from vendor risk swung from 81.5% in H1 '24 to 2.3% this half, a much sharper drop than claim frequency saw over the same period, which fell from 61.3% to 31.8%. When a single vendor event touches dozens of

clients at once, portfolio-level exposure is what matters most. Wildly varying loss figures in different periods reflect the ability for one event to shape the entire year's losses. The lower loss share this half reflects an absence of a wide spread event, not a drop in exposure.


Agentic AI raises the same kind of risk-aggregation scenario at scale: one automated change propagating across every production system at once, with no human in the loop to catch it. This could produce the same shape of failure as the 2024 CrowdStrike outage¹⁵, and a potentially wider blast radius once it happens. Microsoft reports that 80% of Fortune 500 companies use active AI agents¹⁶, and OpenAI claims that 92% of Fortune 500 companies use their products¹⁷ demonstrating the wide adoption and the potential for broad reaching impacts.

Top Causes of Loss


Cause of Loss	2024 % of Claims	2024 % of Losses	2025 % of Claims	2025 % of Losses	H1 2026 % of Claims	H1 2026 % of Losses
Extortion	6.04%	65.26%	6.42%	75.39%	5.80%	73.05%
Transfer Fraud	8.55%	3.56%	10.19%	5.35%	4.71%	9.20%
Wrongful Data Collection	6.04%	3.23%	11.62%	0.00%	16.67%	3.88%
Data Breach	0.00%	0.00%	0.00%	0.00%	5.68%	2.33%
Vendor Data Breach	16.48%	2.15%	19.57%	0.34%	15.10%	1.94%
Vendor - Business Interruption	15.09%	22.15%	15.69%	15.69%	8.45%	0.00%

Cause of Loss Analysis


2024 % of Claims and % of Losses



2025 % of Claims and % of Losses



2026 H1 % of Claims and % of Losses




The ROC connects threat intelligence, data science, claims, and technical engineering into a single function that hunts for risks proactively, before they become losses. Drawing on real loss outcomes: with visibility across a large portfolio of incidents, the ROC can advise clients on how to both prevent and respond to attacks. When a claim does occur, it applies that same integrated knowledge to help clients limit the damage. The anonymized cases below are drawn from recent cases.

The anonymized cases below show what that looks like in practice.

A breach claim that wasn't a breach

A threat actor posted a claim on a criminal forum: they had breached a Resilience client's core customer database and could prove it. The client, a mid-sized online retailer, activated its incident response process and engaged the ROC, and for a moment the claim as posted was the only version of events on the table: a confirmed breach, with everything that follows from one — regulatory notification, legal spend, and a negotiation if a ransom demand came next. Forensic analysis told a different story. There was no evidence of administrative access, no exploitation of vulnerabilities in the internal infrastructure, and no exfiltration from the corporate network or backend databases.

What the threat actor had actually done was credential stuffing against the client's public-facing customer portal — logging in as legitimate customers using credentials already exposed in prior third-party breaches and infostealer dumps — then dressing up what they saw on the portal with previously published leak data to make an old exposure look like a fresh compromise. The ROC classified the claim as unsubstantiated, and the client's disclosure decision followed the forensic findings instead of the threat actor's framing. Telling credential stuffing apart from an actual compromise is what kept a bluff from becoming a notification obligation and a legal bill.



Check Point SmartConsole



New Critical Finding

28 July 2026



High Risk Misconfigured Ports

45

Found
25 Jun 2026

An email flood, then a fake rescue call

A multinational architecture and design firm was hit with an email bomb this half, a wave of inbound email designed to bury an employee in noise. Minutes later, that employee got a Microsoft Teams call from someone claiming to be the internal security team, offering to help clear the mess. The employee installed a remote-access tool at the caller's instruction. From there, attackers deployed credential-harvesting utilities and reached a domain admin account within hours.

The client's EDR solution successfully halted the attack early, limiting the damage to a portion of their systems and a small amount of exfiltrated data. The company determined the criticality of the exfiltrated data was low. The initial ransom demand was \$2M, with the threat actor dropping the demand to \$800K through expert negotiation aided by the ROC. During that time, the client was able to fully restore any encrypted data and chose not to pay for data suppression.

Advanced threat hunting thwarts an attack

In early 2026, the Resilience Risk Operations Center (ROC) identified a sophisticated, emerging threat campaign that utilized specific network management tools in the attack sequence. Utilizing proprietary threat intelligence, the ROC proactively issued a high-priority advisory to a client in the logistics sector, providing specific Indicators of Compromise (IOCs) several days before any breach attempt was launched against the company. Armed with this intelligence, the client's internal security team preemptively hardened their defensive posture, reconfiguring access controls and modifying endpoint detection and response rules to monitor for the exact patterns highlighted by the ROC.

When the threat actor subsequently initiated the breach, they managed to gain an initial foothold within the network.

However, because the client's security systems were already primed to recognize the adversary's signature, the intrusion was detected immediately. The security team intervened decisively while the attacker was still in the early stages of the attack, isolating the compromised system and preventing lateral movement or privilege escalation. This successful near-miss illustrates how shifting the security model from reactive remediation to proactive, intelligence-led defense, Resilience empowered the client to neutralize a sophisticated actor after the initial access was achieved, but well before any operational impact or data loss occurred.

Looking ahead

JADEPUFFER and the Hugging Face breach were warning shots, not a preview of this year's losses. That gap closes the day a criminal points a working agent at profit instead of a benchmark, and that's likely sooner than most H2 plans assume.

Until then, AI's clearest effect on the portfolio is the one already in the data: it makes social engineering more convincing, not more novel. That's why losses tied to a person believing something false have climbed from 17.7% to 85.3% of the portfolio in two years, while known vulnerabilities (7.0%) and governance gaps (5.4%) sit right behind it as fixes that exist but aren't being applied fast or widely enough.

Containment is what will matter on both fronts. How an attack starts, automated or human, matters less than how fast it's caught and how much loss gets limited once it's underway. Building that capability now prepares an organization for whichever kind of attack shows up first.

A few places to start, not a complete list:

- Zero Trust architecture: segmentation, privileged access management, just-in-time access
- Stricter payment controls, especially callback verification and dual sign-off on high-value transfers
- Harder social engineering and phishing tests, not just lower click rates
- Stronger vendor risk management and contingency plans for interruption of critical services.
- More active red teaming, including LLM-aided red teaming

The real task is quantifying value at risk and identifying which of these controls actually reduce it for your organization, then funding those first.

Conclusion

The headlines promised a much bigger AI story than what actually happened. Look at where the losses came from this half, and the answer is still mostly people — someone believed a voice or a message that sounded right, or a policy that existed on paper but never reached every account it was supposed to cover. None of that needed AI to happen.

None of the fixes are exotic, either. Most already exist inside many enterprise security programs: Zero Trust, payment verification, phishing tests calibrated to something harder than a vanity click rate. What's missing is speed and coverage, not invention.

AI does change one real thing: it makes the oldest trick in the book, getting someone to believe something false, more convincing than it used to be. Agentic attacks are a second front worth watching, even without a claim behind them yet.

Both call for a risk-first approach: CISOs, CFOs, and risk managers managing exposure together, not in separate rooms. What separates organizations going into H2 is whether leaders quantify their value at risk before a claim does it for them, then fund the controls that actually move it. Everything else in this report is context for that one decision.

Appendix

Methodology

Data window	All figures reflect Resilience's cyber insurance claims portfolio for the first half of 2026 (January 1 through June 30, 2026) unless noted. Comparisons cover five half-year periods: H1 2024, H2 2024, H1 2025, H2 2025, and H1 2026.
Portfolio scope	Frequency figures are normalized against the average in-force policy count across the period and expressed as claims per 100 policies. That normalization is what accounts for roughly doubled raw claim volume year over year without treating it as risk escalation.
Claim scope	Included are cyber-relevant claims made in the reporting period. Excluded are general-liability and media claims (copyright, false advertising, trademark infringement) that fall outside the cyber definition.
Control-adoption figures	Control-adoption percentages (admin MFA enrollment) are drawn from the subset of clients who self-report security controls, not from the full portfolio. Denominators are the self-reporting population and vary by control.
Cause of loss	The cause-of-loss detail chart shows any cause responsible for more than 3% of incurred losses in any of the five half-year periods. Causes below that threshold are aggregated.
Point of failure	Every claim's point of failure is sorted into one of four categories based on what actually failed: a specific person in a specific moment (Social Engineering), a flaw or gap in the technology itself (Vulnerabilities), a third party (Vendor), or a standing program or policy gap that existed before any attacker showed up (Governance).
Social engineering scope	The 17.7% to 85.3% figure cited throughout the report is a narrower slice than the full Social Engineering category. It includes only phishing, social engineering, and transfer fraud — the sub-types where a threat actor got a person to act on a fraudulent lure. Other sub-types (mismailing, loss of records, insider error variants) are included in the claim count but carried zero incurred loss in every half.
Development	H1 2026 severity figures are shown as reported. The largest H1 2026 cohort is still undeveloped and severity per policy will climb as claims mature.
Payment outcomes	Non-payment and payment figures are described qualitatively in body copy rather than with specific dollar or percentage figures.
Anonymization	All case studies are anonymized. Named public incidents (Canvas, CDK Global, Change Healthcare) are referenced only where public reporting has established the incident. Any related exposure in the Resilience portfolio is described qualitatively rather than by exact figures.
Currency	All figures are expressed in US dollars. Amounts originally in other currencies were converted at flat rates for the period.

Appendix

Endnotes

- ¹ Michael Clark / Sysdig Threat Research Team, "JADEPUFFER: Agentic ransomware for automated database extortion," Sysdig blog, July 2026. Analysis of the first documented case of an end-to-end ransomware operation driven autonomously by a large language model agent. Initial access via CVE-2025-3248 (Langflow RCE), pivoting to a production MySQL / Alibaba Nacos configuration server; 1,342 Nacos configuration items encrypted with an unrecoverable key. Behavioral evidence for LLM-driven operation includes natural-language reasoning in payloads, machine-speed self-correction (failed login to working fix in 31 seconds), and 600+ distinct purposeful payloads in a compressed window. sysdig.com/blog/jadepuffer-agentic-ransomware-for-automated-database-extortion
- ² OpenAI, "OpenAI and Hugging Face partner to address security incident during model evaluation," July 21, 2026. openai.com/index/hugging-face-model-evaluation-security-incident
- ³ World Economic Forum (in collaboration with Accenture), "Global Cybersecurity Outlook 2026," January 12, 2026. Fifth-edition annual survey drawing on responses from 804 qualified participants across 92 countries, including 105 CEOs and 316 CISOs. In 2026, cyber-enabled fraud and phishing displaced ransomware as the top-ranked cyber concern among surveyed CEOs, with AI-related vulnerabilities newly ranked second. weforum.org/publications/global-cybersecurity-outlook-2026
- ⁴ Verizon, "2026 Data Breach Investigations Report," May 2026. 19th annual edition, based on analysis of security incidents from November 2024 through October 2025. Found that 73% of ransomware victims had an infostealer or credential-leak event in the year before the attack, and half of those had that event within 95 days of the ransomware deployment. verizon.com/dbir
- ⁵ Troutman Pepper, "SB 690 Amended: California Moves to Strip Private Right of Action for Pen Register Claims," July 2, 2026. troutmanprivacy.com/2026/07/sb-690-amended-california-moves-to-strip-private-right-of-action-for-pen-register-claims-2
- ⁶ Mandiant (Google Cloud), "M-Trends 2026," May 2026. Based on more than 500,000 hours of incident response investigations conducted in 2025. Found mean time-to-exploit has gone negative, to an estimated -7 days — exploitation now often begins before a patch is available. cloud.google.com/blog/topics/threat-intelligence/m-trends-2026
- ⁷ CrowdStrike, "2026 Global Threat Report," February 2026. Found average eCrime breakout time (the interval between initial compromise and first lateral movement) of 29 minutes, down from 48 minutes in 2024; fastest observed breakout was 27 seconds. In one documented intrusion, data exfiltration began within four minutes of initial access. crowdstrike.com/en-us/global-threat-report/
- ⁸ Check Point Research, "The State of Ransomware — Q1 2026," May 2026. Analysis of more than 70 active data leak sites listing 2,122 new victims in Q1 2026. The top ten named ransomware groups accounted for 71.1% of victim postings, up from 57% in Q3 2025. research.checkpoint.com/2026/the-state-of-ransomware-q1-2026
- ⁹ Jud Dressler, "The Ransom Dilemma: Driving Loss Control in the Age of Extortion-Only Attacks," Resilience whitepaper, 2026. Based on Resilience's claims portfolio: data-theft extortion without encryption grew from 49% of extortion claims in H1 2025 to 65% in H2 2025, with pure-encryption attacks falling to 13% by year-end. unlock.cyberresilience.com/the-ransom-dilemma
- ¹⁰ Google Threat Intelligence Group (GTIG), "2026 Ransomware Landscape Report," March 2026 (as reported by CyberScoop, March 16, 2026). Data-theft extortion without encryption grew from approximately 2% of directly observed financially motivated incidents in 2020 to more than 15% in 2025. cyberscoop.com/google-threat-intelligence-group-ransomware-report-2026
- ¹¹ CrowdStrike, "2025 State of Ransomware Survey," October 21, 2025. Survey of 1,100 global security leaders. Found that 83% of organizations that paid a ransom were attacked again, and 93% had data stolen regardless of payment. crowdstrike.com/en-us/resources/reports/state-of-ransomware-survey/
- ¹² Jessica Lyons, "Double Canvas intrusion confirmed as ShinyHunters resets leak deadline," The Register, May 12, 2026. ShinyHunters claimed approximately 3.65 TB of data including approximately 275 million records from approximately 8,800 institutions. theregister.com/security/2026/05/12/double-canvas-intrusion-confirmed-as-shinyhunters-resets-leak-deadline/5238361
- ¹³ U.S. Securities and Exchange Commission, UnitedHealth Group Form 8-K, February 22, 2024. Ransomware attack on Change Healthcare, a UnitedHealth Group subsidiary, disrupted claims processing and eligibility verification across the U.S. healthcare system beginning February 21, 2024. sec.gov/Archives/edgar/data/731766/000073176624000045/unh-20240221.htm
- ¹⁴ Bloomberg / Reuters, June 2024 coverage of the CDK Global cyberattack. Ransomware attack on June 18–19, 2024 attributed to the BlackSuit ransomware group. Disrupted dealer management systems used by approximately 15,000 U.S. and Canadian auto dealerships for roughly two weeks. bloomberg.com/news/articles/2024-06-24/blacksuit-cybercrime-gang-blamed-in-major-cdk-ransomware-attack
- ¹⁵ CISA, "Widespread IT Outage Due to CrowdStrike Update," July 19–20, 2024. A faulty content update to CrowdStrike's Falcon sensor caused approximately 8.5 million Microsoft Windows systems to crash worldwide, disrupting aviation, healthcare, banking, and emergency services. cisa.gov/news-events/alerts/2024/07/19/widespread-it-outage-due-crowdstrike-update
- ¹⁶ Microsoft Security, "80% of Fortune 500 Use Active AI Agents: Observability, Governance, and Security Shape the New Frontier," February 10, 2026. Microsoft's Cyber Pulse report found that more than 80% of Fortune 500 companies had an active AI agent built with Microsoft Copilot Studio or Microsoft Agent Builder, based on first-party telemetry measuring agent activity during the 28 days ending November 2025. <https://www.microsoft.com/en-us/security/blog/2026/02/10/80-of-fortune-500-use-active-ai-agents-observability-governance-and-security-shape-the-new-frontier/>
- ¹⁷ VentureBeat, "OpenAI Says ChatGPT Now Has 200M Users," August 29, 2024. OpenAI confirmed to VentureBeat that ChatGPT had surpassed 200 million weekly active users — double the figure from ten months earlier — and disclosed that 92% of Fortune 500 companies were using its products, coinciding with ChatGPT Enterprise's first anniversary. <https://venturebeat.com/ai/openai-says-chatgpt-now-has-200m-users>

See the latest trends and
analysis in cyber claims

