veeam

coveware by veeam

# Inside the Economics of Cyber Extortion

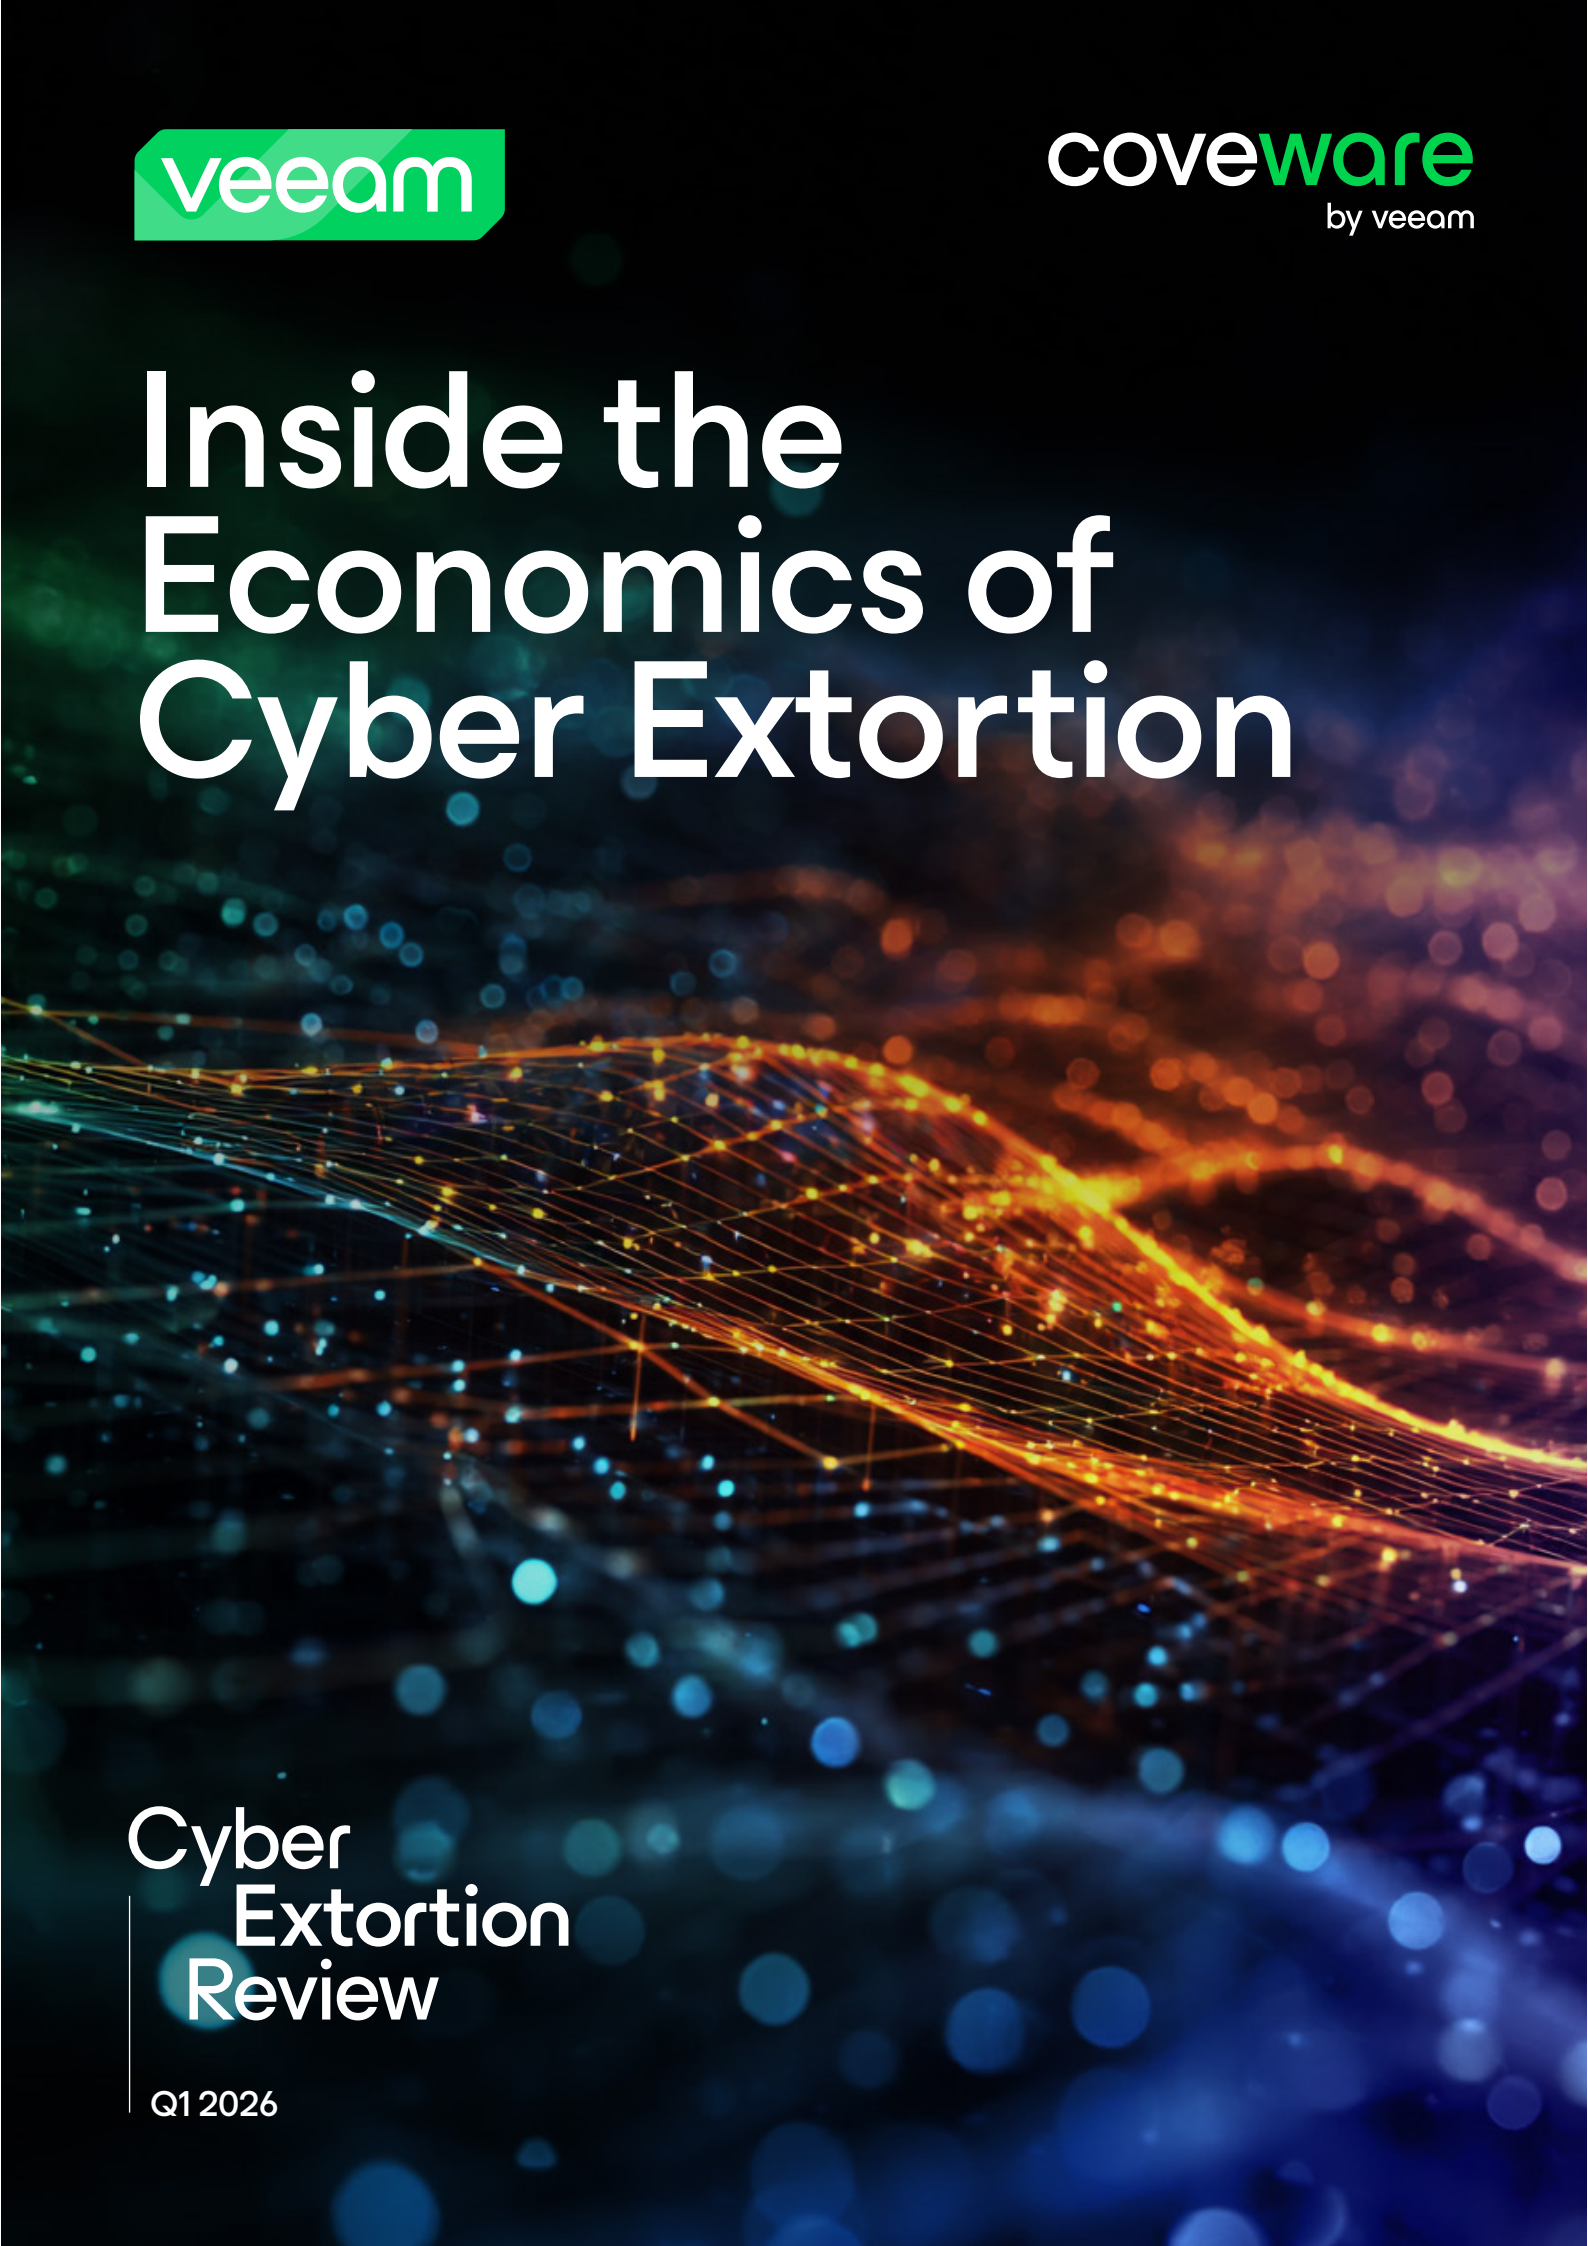

## Cyber Extortion Review

Q1 2026

Cyber extortion is moving faster as AI reshapes vulnerability discovery. Attackers are compressing response windows, forcing executives to make sharper decisions before disruption becomes leverage.

# Foreword

The past decade has seen cyber extortion evolve from a technical security issue into a material business risk for all executives. The conditions driving extortion are changing rapidly, as AI-assisted vulnerability discovery and exploit development, shrinking patch windows, and faster exploitation cycles reduce the time organizations have to respond before attackers operationalize and exploit those weaknesses.

Q1 saw the arrival of Mythos AI, a meaningful signal of where vulnerability discovery may be heading. Anthropic's model autonomously identifies software flaws and can generate exploits. The shift from human to AI detection means more flaws surfacing, and sharper pressure to patch them fast.

In this report, we examine Coveware by Veeam's Q1 2026 data and generate insights for executives to reduce attacker leverage, strengthen resilience, and improve decision-making before, during, and after an incident.

*Coveware by Veeam's data reflects observed casework, not a representative market sample.*

**"Faster exploitation cycles reduce the time for organizations to respond"**

# A Quarter in Review: Q1 Data at a Glance

Average ransom payments rose 15% to $680k in Q1, while median payments fell 7% to $300k, consistent with a mix shift toward a smaller number of high-leverage, high-payment cases. Critical sectors such as the public sector and healthcare accounted for 29% of attacks, reinforcing that business interruption remains a stronger payment driver than data exposure alone.

## Essential services are in the crosshairs

**29%**

of attacks targeted organizations that cannot afford to go dark such as hospitals, schools and public services

## All ransomware payment resolution rates, Q1 2024–Q1 2026

Extortion pressure is converting less than in previous years, not more

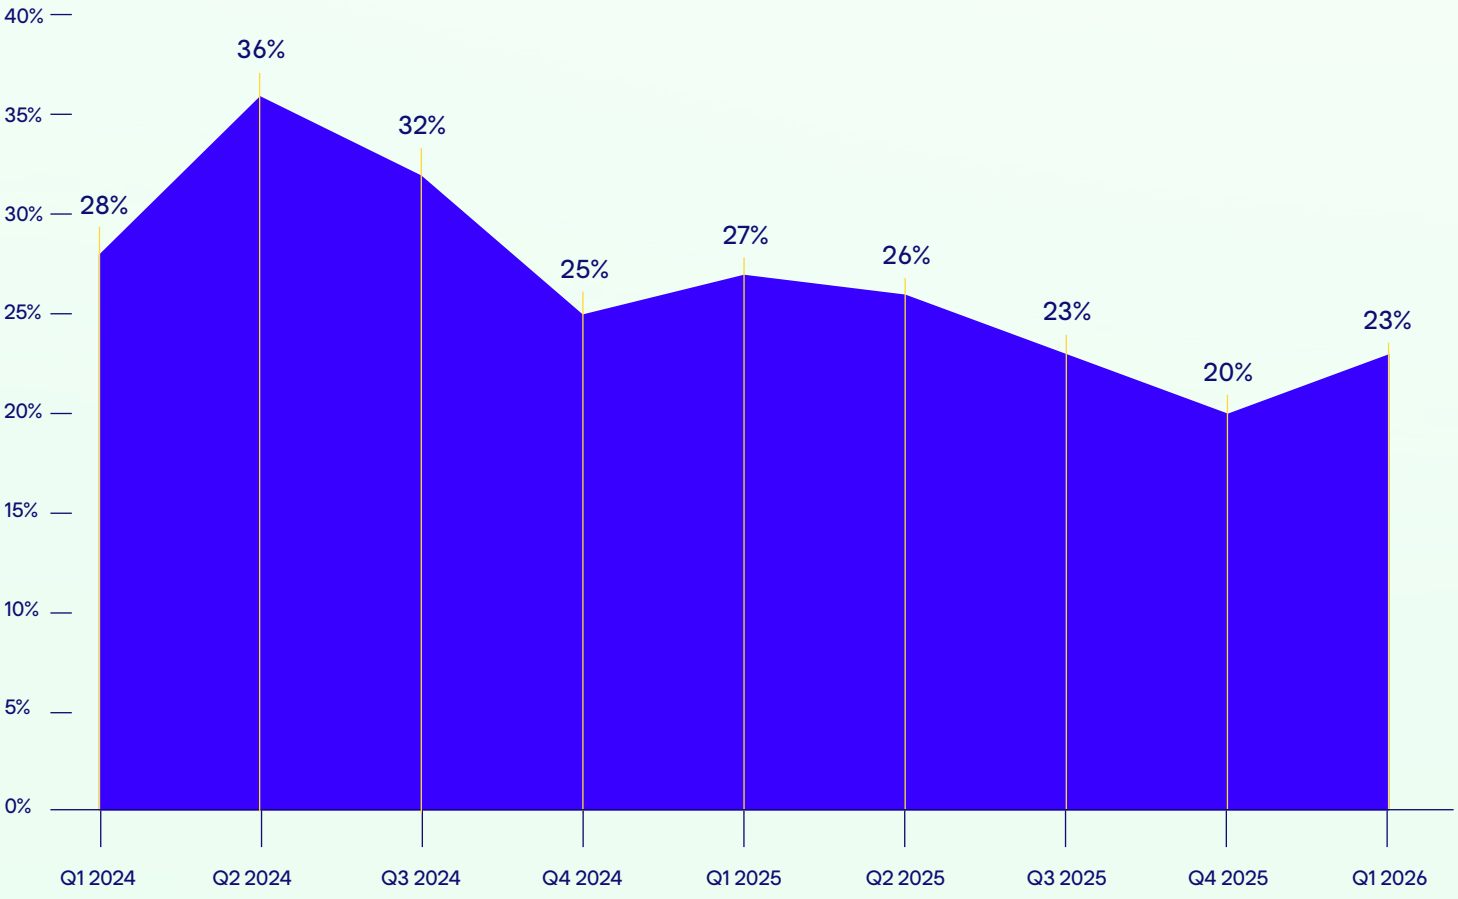

## Most common ransomware variants, Q1 2026

The market remains split between encryption-led and data-theft-only operators

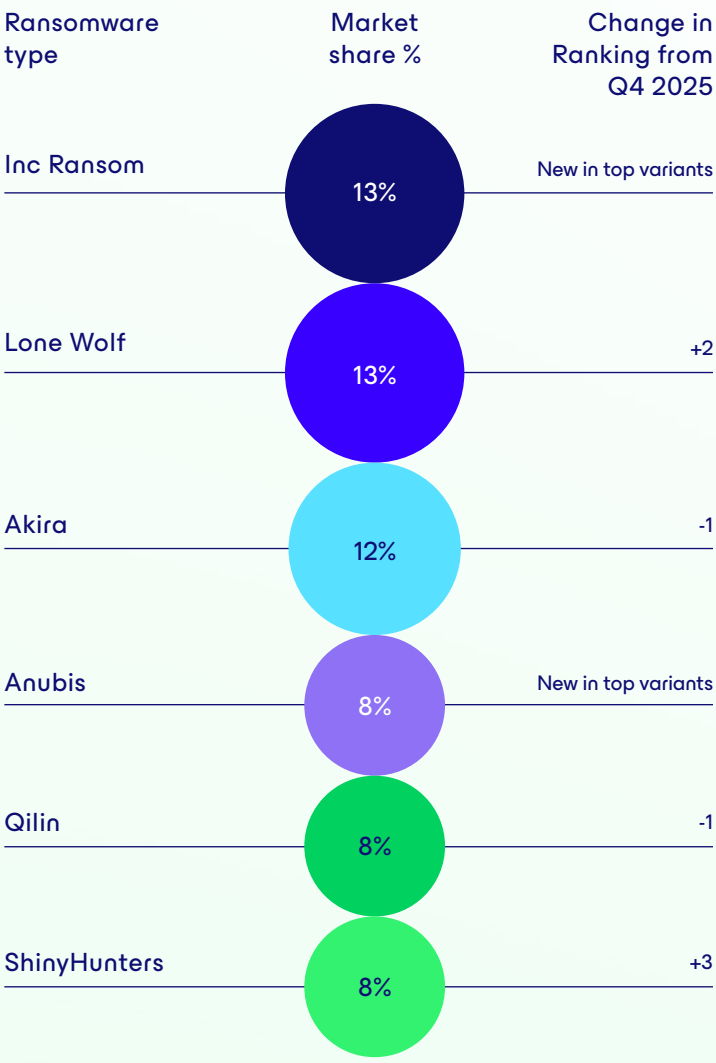


## Ransomware-impacted companies by employee size, Q1 2026

Extortion still lands hardest in the mid-market

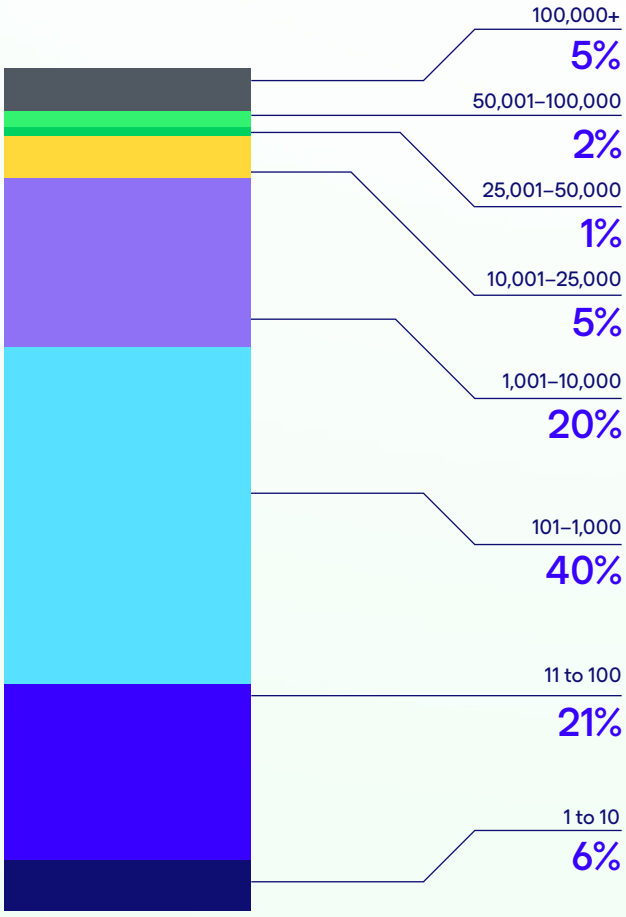


## Industries impacted by ransomware, Q1 2026

Healthcare and consumer services are the most targeted sectors for ransomware

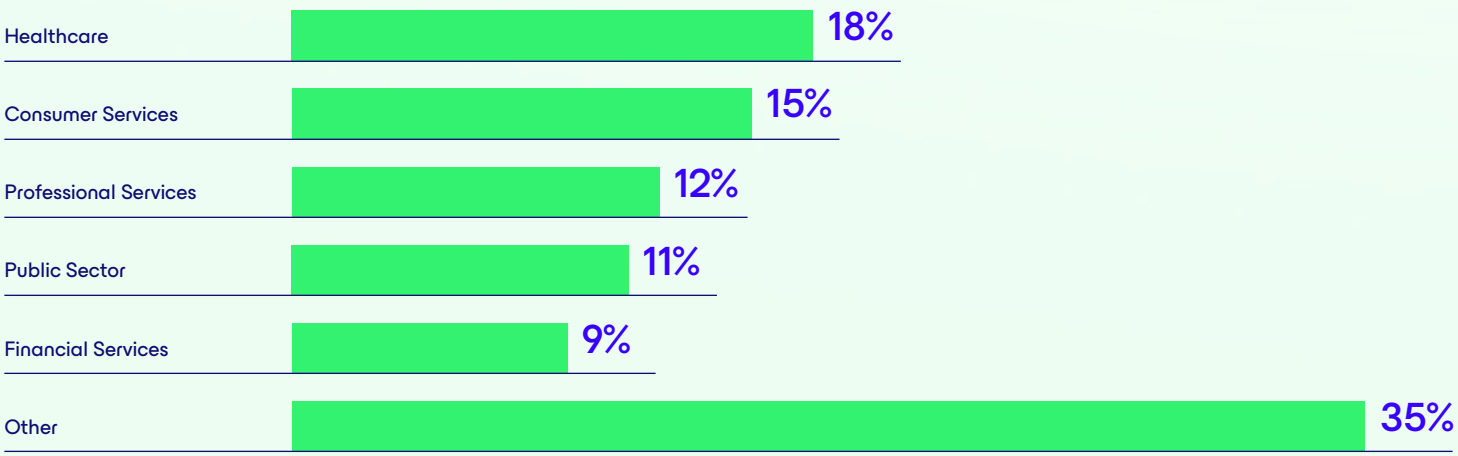

When machines find flaws faster than teams can respond, every patch decision becomes a negotiation between speed, confidence and business continuity

ARTICLE ONE

# From the Negotiation Room

*AI-discovered vulnerabilities leave security teams less time to decide how to respond. What follows is the perspective of a typical enterprise CISO.*

When a new vulnerability surfaces, my team's response is well rehearsed. We confirm whether the flaw actually affects the business, test the patch, and roll it out without taking servers and users offline. Done well, that process lets me, the CISO, bring leadership an uncommon piece of good news: "We're not vulnerable to this."

Increasingly, though, the discovery itself comes from a different source. A flaw that once would have reached us through a researcher, a monthly update, or a vendor's out-of-band advisory now arrives from AI. Anthropic has given Mythos the ability to find software flaws and generate exploits, with early signals that it may even surface major zero-days. A patch may well exist, but the harder question now is how seriously my team should treat what an autonomous system has produced.

My team has always been wary of automated vulnerability discovery, and the doubts are reasonable. We question whether a flagged flaw is genuine, whether a machine can judge severity accurately, and whether it can assign a CVSS score reliable enough to drive patch prioritization. Those doubts become more pressing with Mythos, which produces findings faster than we can vet them.

This marks a real shift in vulnerability response. The old question was whether the patch existed; the new one is how quickly we can make a decision we can defend. As machines surface more flaws, and surface them faster, the bottleneck moves from engineering to judgment. We must weigh how long the business can stay exposed, how fast it can move without breaking things, and whether risk-based patching still holds once AI has compressed the timeline it was built on.

Mythos can identify a flaw, but determining who is likely to exploit it is a separate question, and the second decision we face alongside the first.

**"We must weigh how long the business can stay exposed, how fast it can move without breaking things, and whether risk-based patching still holds"**

## 48 hours

The highest-risk classes need even faster patching deployment speeds to match attacks that are increasing in speed and autonomy

When we know the actor behind a threat and understand their TTPs, we can anticipate their likely capabilities, preferred exploit paths, and operational tempo, and those factors determine how urgently we respond. Attribution, however, carries its own uncertainty. It is especially difficult with lone operators, where we may lack the intelligence to link events to a known actor. Acting on the wrong attribution creates its own problems: we may spend resources against a threat that cannot act, or underestimate an adversary more capable than we expect.

We now make both decisions, the patch and the attribution, faster and with less information than we would like. The outcome we hope for, "We're not vulnerable to this," remains possible, but it is becoming rarer and harder to state with confidence. Making sound, defensible decisions within the shorter timelines that AI-driven discovery imposes is the central challenge I now face.

## Who's Most Exposed and Why

Our Q1 2026 data shows that healthcare accounted for 18% of ransomware incidents, consumer services 15%, professional services 12%. Financial services remained an attractive target for the usual reasons: concentration of sensitive data and financial leverage during negotiations. Sector is only part of the picture. Three structural risk profiles recur across incidents regardless of industry.

The first is legacy technical debt, particularly internet-facing systems that can't be patched quickly without disrupting operations. Many of these environments run on platforms once considered stable but now exposed to automated exploitation.

The second is organizations building out AI infrastructure. These environments introduce heterogeneous hardware such as GPUs, ARM hosts, embedded firmware, and externally managed components. Operational complexity climbs, and patching requires downtime that businesses can't absorb.

The third is modern development practice. AI-assisted coding and accelerated delivery are introducing unmapped dependencies, externally sourced components, and AI-generated code into production environments faster than governance and security can track.

In all three cases, the core issue is visibility and control. Organizations lose leverage when they can't say what's running, who maintains it, how fast it can be patched, or which dependencies could fail during an incident. That uncertainty is what attackers exploit.

### Attack vector distribution, Q1 2026

Access and vulnerability gaps dominate Q1 incidents

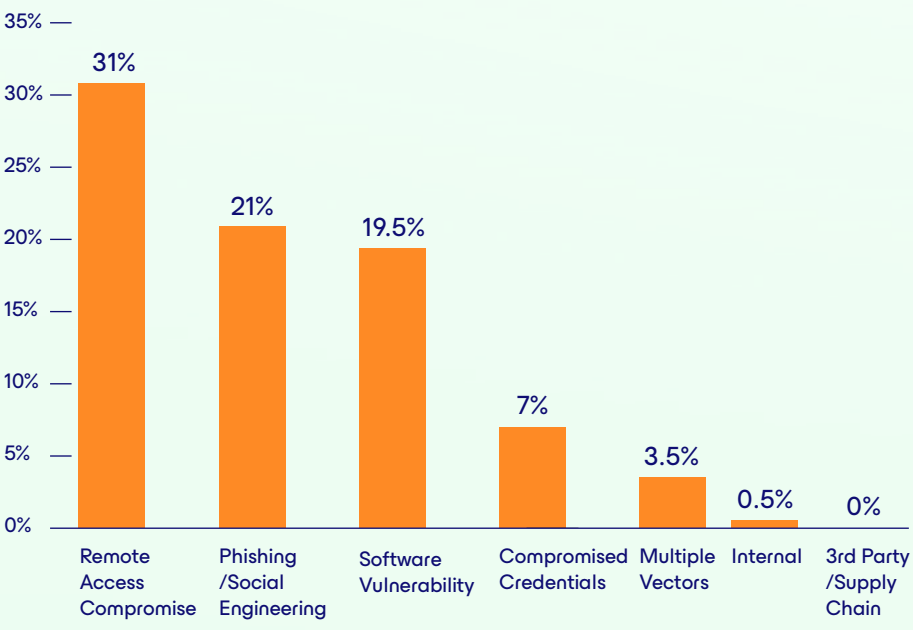


### Legacy code is no longer safe

**27 years**

The Mythos model identified a 27-year-old vulnerability in code that had passed extensive automated testing

### The cost of discovery

**$20,000**

The estimated operational cost of finding the OpenBSD vulnerability

ARTICLE TWO

# The Identity Perimeter

Attackers are no longer just breaking in through technical flaws. They are logging in through trusted identities, turning access, visibility and response speed into the new front line of cyber extortion.

Review any routine patch cycle and the same pattern emerges: most critical fixes target remote code execution (RCE), the vulnerabilities that let attackers execute malicious code on a system without requiring physical access or user interaction.

Remote access is now essential because of hybrid work, distributed infrastructure, SaaS adoption, outsourced operations, and the need for constant connectivity. From an attacker's perspective, those same pathways act as a front door. RCE vulnerabilities give them a foothold on internet-facing systems, from which they can drop malware, steal credentials, or establish persistence.

Focusing purely on exploitation risks, however, misses the larger shift taking place across modern intrusions. Attackers are increasingly succeeding through identity compromise, not break-ins. They show up as legitimate users because they've authenticated as legitimate users. That changes both detection and response from the ground up. Often you don't even know there's been an attack, let alone what action needs to be taken.

"Remote access compromise" has become a term too broad to be useful. In many incidents, attackers aren't exploiting remote access tools directly; rather the perimeter has moved from the network edge to the identity layer. Mechanisms such as SSO, MFA recovery workflows, SaaS integrations, and OAuth grants are now the attack surface that matters. Once attackers have trusted access through identity systems, traditional security controls become far less effective at distinguishing malicious activity from legitimate behavior.

Coveware by Veeam incident data reinforces this reality. Lateral movement occurs in 79% of cases, data exfiltration in 73%, and measurable operational disruption in 58%. Each of these phases creates an opportunity for detection, but only if organizations maintain visibility into identity-layer activity.

## “Prioritize remediation around the systems that create the greatest leverage for attackers”

Ray Umerley, Field CISO for Coveware by Veeam, said: “These are not opportunistic smash-and-grab attacks. They are structured campaigns involving reconnaissance, privilege escalation, persistence, and staged execution.”

Phishing needs to be similarly reconsidered; it is increasingly just a delivery mechanism. “By the time an intrusion presents as remote access compromise, the compromise may already be well established,” Umerley added. “This is a test for visibility and your resilience.”

At the same time, patching dynamics continue to accelerate. Once vendors release security updates, attackers can rapidly reverse-engineer patches and automate exploitation. In environments where deployment cycles still operate in days rather than hours, defenders are often reacting after the exposure window has already opened.

As a result, organizations need to prioritize remediation around the systems that create the greatest leverage for attackers. Internet-facing infrastructure, identity providers, remote management platforms, externally exposed SaaS integrations, and third-party access pathways all require additional attention and protection.

Modern extortion resilience increasingly depends on an organization’s ability to reduce attacker dwell time through faster visibility, tighter identity controls, and more mature response processes.

### Top observed tactics, Q1 2026

Most extortion events are hands-on operations with multiple points of intervention

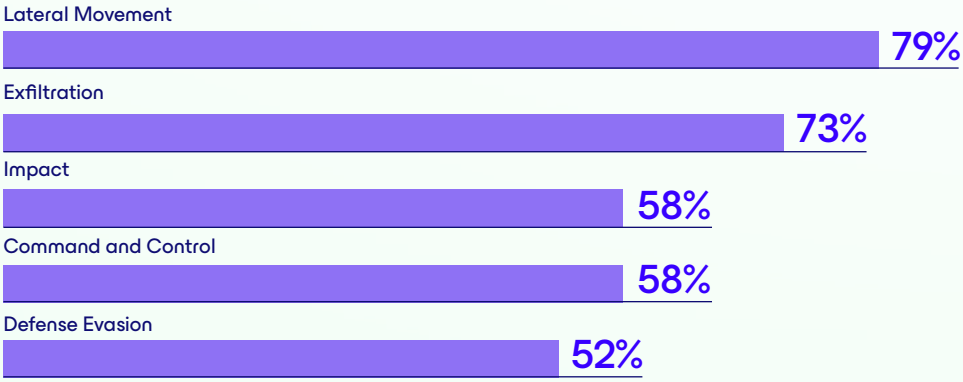


### Four Signals That the Identity Perimeter Is Already Under Pressure

Identity-based intrusions rarely begin with ransomware deployment. Most campaigns generate warning signals well before complex disruption occurs. The challenge is recognizing them early enough to act decisively.

**Anomalous MFA resets or new enrollments.** Attackers frequently attempt to register new authentication methods after obtaining valid credentials. Sudden MFA resets, unexpected device enrollments, or repeated authentication recovery attempts should be treated as high-confidence indicators of compromise.

**New OAuth grants or connected application permissions.** OAuth abuse allows attackers to maintain persistent access without repeatedly triggering password resets or MFA challenges. Unfamiliar third-party application permissions, especially involving email, file storage, or collaboration platforms, require immediate scrutiny.

**Unusual help desk or identity recovery activity.** Social engineering against internal support teams remains highly effective. Attackers impersonate employees to request password resets, device re-enrollment, or MFA bypasses. Identity recovery workflows increasingly represent a frontline security control rather than an administrative process.

**Privileged access from unfamiliar locations or unmanaged devices.** Logins from unusual geographies, unmanaged hardware, or previously unseen devices remain strong indicators of account takeover activity. In many incidents, these events signal that attackers already possess valid credentials and are beginning lateral movement or privilege escalation.

Attackers are becoming more selective, targeting organizations where disruption creates the greatest pressure to pay. Reducing extortion risk now means shrinking the leverage attackers can create once inside.

ARTICLE THREE

# The Leverage Economy

Modern cyber extortion is run on what's worth calling a "leverage economy," or the ability to amplify outcomes beyond direct effort through capital, systems, or structural advantage. An attacker's aim is to invest as little effort as possible while creating disproportionately large financial and downtime consequences for victims.

Often more effort goes into an attacker's selection process than the attack itself, as modern ransomware groups carefully consider the price of disruption. Every intrusion is a calculation built around business disruption, recovery timelines, regulatory exposure, reputational pressure, and executive decision-making dynamics. Attackers select victims based on how these factors play into their methodology. They also seek to curate an environment in which the theft and encryption of information must be met with ransom payment — at least, to make payment seem like the only commercially rational option.

This is also why effective business disruption campaigns drive the highest payments. Attackers maximize the weight of their punch by targeting those they will hurt the most, knowing that pressure is not equal among all organizations. Leverage grows with the gap between who experiences the pain of downtime and who authorizes payment. In healthcare, disruption affects patients. In education, students and staff. In consumer services, customers and revenue. The wider this gap becomes, the greater the leverage attackers gain during negotiations.

Coveware by Veeam's incident data confirms this. "Must-operate" sectors are among the most heavily targeted because disruption rapidly escalates into executive and reputational pressure. This is particularly visible across mid-market organizations. Many operate critical services without the resilience, segmentation, or recovery maturity typically available to larger enterprises, and attackers understand this imbalance. Low tolerance and constrained recovery options create strong payment pressure.

In other campaigns, traditional ransomware encryption groups hit a target, created disruption, and stole data. "But they're very focused on the impact on that organization and trying to pursue monetary gain as a result," said Ray Umerley, Field CISO at Coveware by Veeam.

The extortion market itself is also becoming more economically segmented. Coveware by Veeam's Q1 data showed average payments up 15% to $680,000, while median payments fell 7% to $300,000.

### The market has found its sweet spot

**66%**

of all Q1 attacks targeted organizations with between 11 and 1,000 employees

### The median victim is getting bigger

**+150%**

The median victim size jumped to 500 employees in Q1 2026: up 150% on Q4 2025

### Ransom payment amounts, Q1 2026 vs Q4 2025

Average payments are rising even as fewer organizations pay

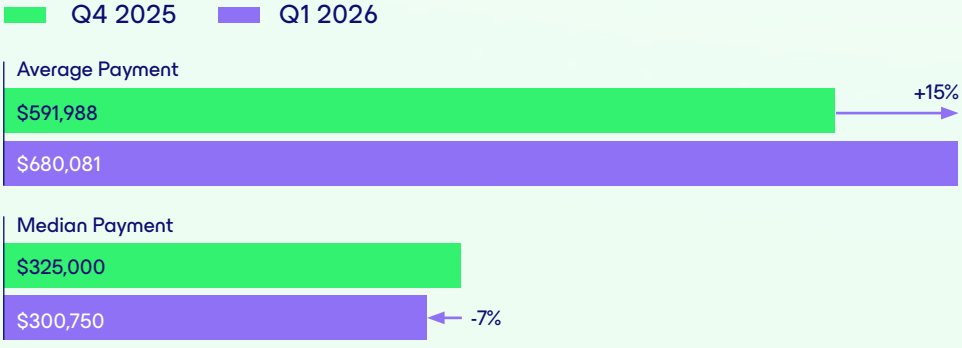

> **"Organizations are far more likely to pay when core operations stop functioning than when facing reputational exposure alone. As such, business interruption remains the primary driver of extortion economics."**

This divergence matters, as it suggests a smaller number of large enterprises are making much larger payments. However, Umerley said "a broader portion of mid-market victims are either recovering independently or refusing to pay lower-leverage demands."

This is not evidence that more organizations are paying overall. Instead, it reflects how attackers are concentrating on the incidents that generate maximum leverage.

Encryption-led groups like Akira and Qilin drive most payments because they create company-wide paralysis. By contrast, data-exfiltration-only actors, including ShinyHunters and Lone Wolf operators, often struggle to generate equivalent payment pressure without disruption tactics.

Organizations are far more likely to pay when core operations stop functioning than when facing reputational exposure alone. As such, business interruption remains the primary driver of extortion economics.

Attackers understand which organizations can absorb reputational damage, restore from backups, withstand public disclosure, or tolerate temporary exposure. As a result, many are abandoning lower-yield extortion models in favor of campaigns designed to maximize disruption quickly and efficiently.

The leverage economy is becoming more selective, more commercially rational, and more asymmetric. As such, defenders now carry two jobs to reduce extortion risk: keep attackers out, but shrink what they can do once they're in.

## Most common ransomware variants grouped by extortion model, Q1 2026

Encryption-led groups continue to drive the bulk of actual payments

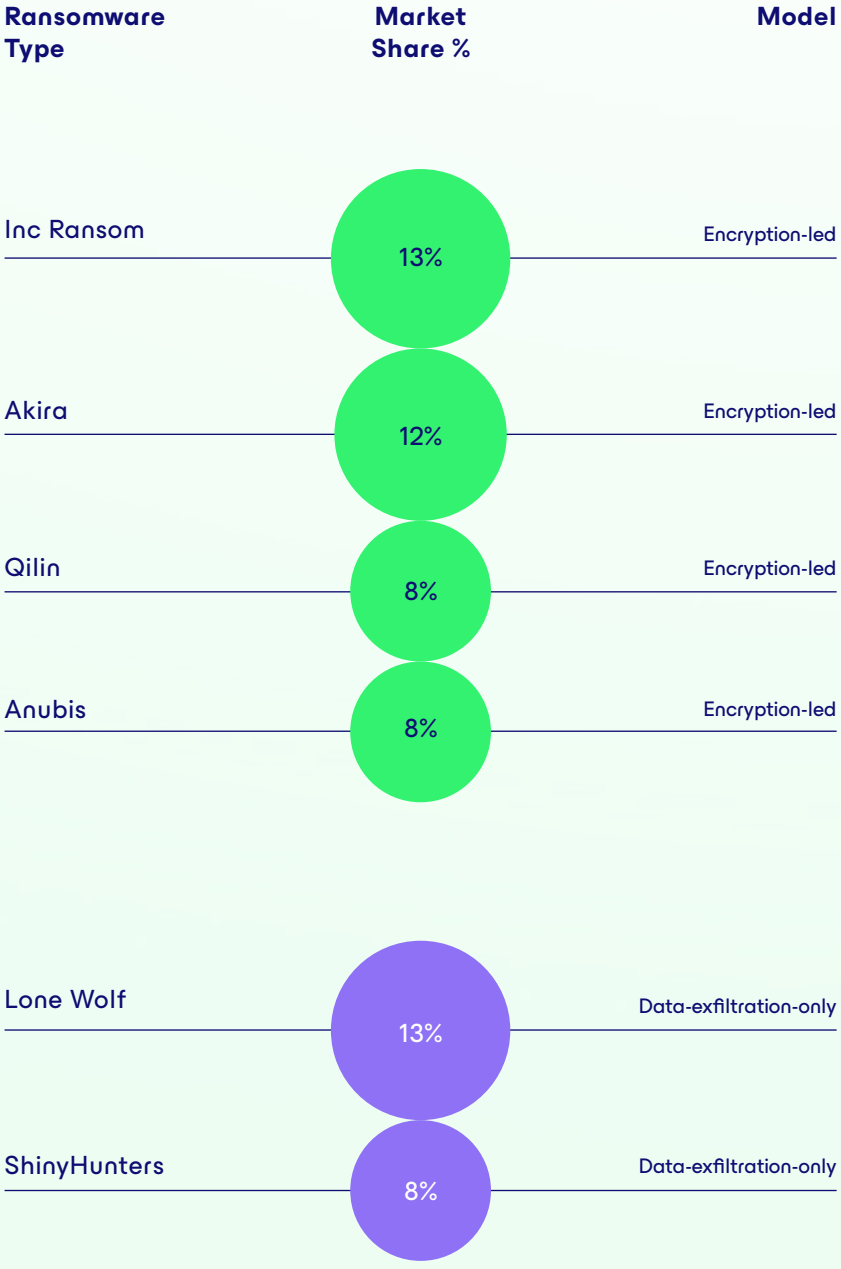

# Conclusion

The data from Coveware by Veeam shows that only one in four supported victims are choosing to make a ransom payment. This is driven by their stronger recovery capabilities, better backups, earlier engagement with third-party experts, and more disciplined executive decision-making.

The willingness to make a payment has declined even as extortion becomes more calculated, with attackers increasingly focusing on organizations where downtime, regulatory exposure, or reputational damage can be leveraged to cause maximum harm quickly.

Organizations affected by these attacks are refusing to buckle under pressure. Non-payment remains the default position, while attackers relying on data-exfiltration-only tactics are seeing those efforts become less effective.

## Non-payment is becoming structural

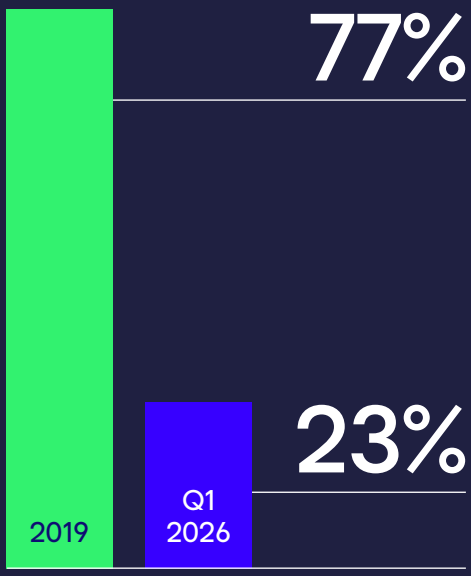


Payment rates have fallen from 77% in 2019 to 23% in Q1 2026

"Good news; the leverage appears to be shifting toward defenders in more cases," said Ray Umerley, Field CISO for Coveware by Veeam. "Consider that payment rarely ends the exposure: data deletion promises are unverifiable, re-extortion can materialize months later, and the short-term appearance of promises kept should not be mistaken for durable risk reduction."

## Preparing for Q2

The emergence of Mythos and its ability to identify new vulnerabilities points to a new trajectory in both AI capability and vulnerability proliferation, potentially driving down the price of exploits.

Organizations that act on this quarter's data will be better positioned than those that wait, and fewer will face the decision of whether to make a cyber extortion payment.

Veeam is the Data and AI Trust Company, specializing in helping organizations ensure their data and AI are fully understood, secured, and resilient to enable the acceleration of safe AI at scale. As the market leader in both data resilience and data security posture management, Veeam is built for the convergence of identity, data, security, and AI risk.

Veeam delivers deep contextual intelligence across every data asset, identity, and AI model. The company governs access for both humans and AI agents, automates privacy, compliance, and remediation processes, and protects and recovers organizations from modern threats – including ransomware, disasters, AI errors, and ensuring the restoration of clean, trusted data. Veeam empowers organizations to move beyond simply protecting data, enabling them to activate and unlock its full potential.

Headquartered in Seattle with offices in more than 30 countries, Veeam protects over 550,000 customers worldwide, including 82% of the Fortune 500, who trust Veeam to keep their businesses running. Learn more at www.veeam.com or follow Veeam on LinkedIn @veeam-software and X @veeam.