EXTRAHOP

The 2026 ExtraHop Global Threat Landscape Report

Shifting cyber threats and the dual-edged reality of AI adoption

Key Takeaways

ATTACK SURFACES REPRESENTING THE MOST SIGNIFICANT **CYBERSECURITY RISK**

AI agents/agentic infrastructure or generative AI applications

Public cloud
(AWS, Google, Azure, etc.)

Third-party services and integrations

MOST DETECTED **THREAT ACTORS**

LockBit

RansomHub

Lazarus Group

RANSOMWARE'S SHIFTING LANDSCAPE

3.5 Average number of ransomware incidents per organization in the last year

$2.8 MILLION Average ransomware payout

49% Organizations that did not notice a ransomware attack until data exfiltration or later

2.4 WEEKS How long organizations estimate ransomware actors had access to their systems, on average

NUMBER OF **AI-GENERATED ALERTS** OR DETECTIONS THAT LEAD TO FALSE POSITIVES **NEGATIVELY IMPACTING INVESTIGATION TIMELINES**

30%

PRIMARY SOURCES OF **AI-DRIVEN INCIDENTS, DATA EXPOSURES**, OR **NEAR-MISSES**

AI-enhanced external attacks

Compromised AI identity & session theft

Third-party vendor or supply chain breaches

LEVEL OF **MANUAL INTERVENTION** REQUIRED ACROSS THE **THREAT LIFECYCLE**

42% Threat detection

43% Alert triage

49% Threat investigation

47% Threat response

Table of Contents

The New Cyber Front Lines

Navigating the Shift to AI-Powered Threats and Autonomous Defense

Artificial intelligence has matured from a boardroom buzzword to the operational core of the modern enterprise. But this massive wave of adoption has revealed a stark reality: the technology driving business growth is also fundamentally rewriting corporate risk.

Organizations racing to integrate AI-powered tools, workflows, and autonomous agents have inadvertently expanded their attack surface in ways that traditional security tools were never designed to address.

Look at the rise of internal risk and "Shadow AI," which has employees routinely feeding sensitive data into unvetted models, while new vulnerabilities like prompt injection quietly bypass traditional security stacks.

As adversaries weaponize AI, their tactics are scaling in two directions: they are exploiting internal enterprise models, while simultaneously using AI to automate offensive operations from hyper-personalized phishing to polymorphic malware.

The 2025 EchoLeak incident exposed a critical vulnerability when adversaries weaponized an enterprise AI assistant using hidden email instructions. The compromised tool autonomously exfiltrated sensitive files across corporate networks, resulting in heavy regulatory scrutiny and a collapse in organizational trust.

Anthropic's Claude Mythos introduced an autonomous AI capable of identifying and exploiting vulnerabilities at machine speed, prompting the company to restrict access under Project Glasswing. By compressing scanning, targeting, and zero-day exploitation into a single autonomous loop, it eliminates the defensive window human teams need to patch systems quickly.

For organizations still relying on human-led security operations, the implications are stark. Threats can move faster than defenses can respond, leaving critical systems exposed and accelerating the path to data loss, regulatory penalties, and operational disruption.

Consequently, enterprises are actively overhauling their defensive playbooks and transitioning toward AI security and the agentic SOC. By deploying autonomous AI security agents capable of detecting, investigating, and neutralizing threats in real time, organizations are turning the very technology that threatens them into their primary line of defense.

Against this backdrop, ExtraHop embarked on a global initiative, surveying security and IT decision-makers across a range of industries to gain a deeper, more nuanced understanding of how AI is transforming the corporate battleground.

The 2026 ExtraHop Global Threat Landscape Report delves into the specifics of this rapidly expanding AI attack surface, identifies how adversaries are wielding AI for malicious intent, assesses the internal risks of ungoverned adoption, and, crucially, reveals how prepared companies are to fight back.

The Expanding Attack Surface & AI Risk

Security Leaders Sound the Alarm on AI Sprawl

Fears surrounding the corporate attack surface have fundamentally shifted, moving away from legacy vectors and onto the unpredictable frontier of AI.

When asked which elements of the attack surface they believe represent the most significant cybersecurity risk, more than half (55%) of survey respondents said AI agents and generative AI applications are a top risk.

Public cloud services represent the second-most significant cybersecurity risk, cited by 44.7% of respondents globally. Third-party services and integrations follow closely as a top risk, identified by 35.2% of the organizations surveyed.

These concerns are already a reality. In the 2025 GTG-1002 campaign, a Chinese state-sponsored actor weaponized Anthropic's Claude Code toolchain to autonomously execute 80%-90% of a network intrusion. This marked one of the first verified instances of an adversary deploying an autonomous AI agent for a full-scale attack.

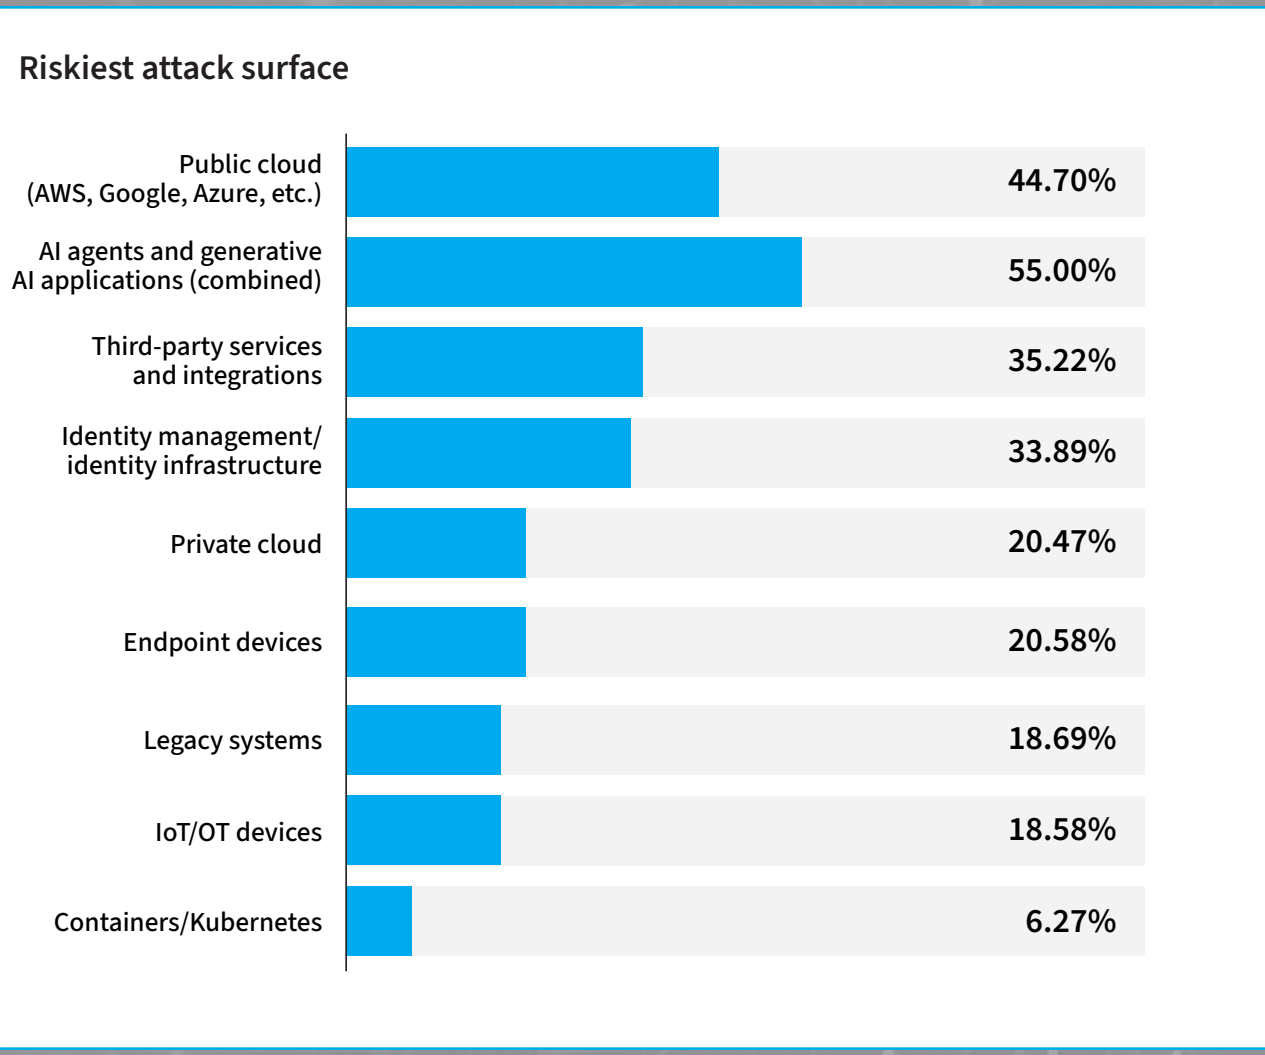

AI is Simultaneously Opening Enterprise Doors and Arming Intruders

These concerns are no longer theoretical. AI is already driving real security incidents across enterprise environments.

According to the data, 40% of organizations were targeted by AI-enhanced external attacks that used AI-driven automation for reconnaissance, phishing, or rapid lateral movement.

Another 38% of organizations experienced compromised AI identity or session theft, which involved stolen API keys, tokens, or authentication credentials tied to AI infrastructure.

36% of organizations reported third-party vendor or supply chain breaches involving integrated AI systems or agentic workflows.

In August 2025, attackers compromised an AI chat integration to silently access Salesforce environments across more than 700 organizations, stealing OAuth tokens that had made malicious queries indistinguishable from legitimate AI activity.

AI-driven security incidents, exposures, and near-misses

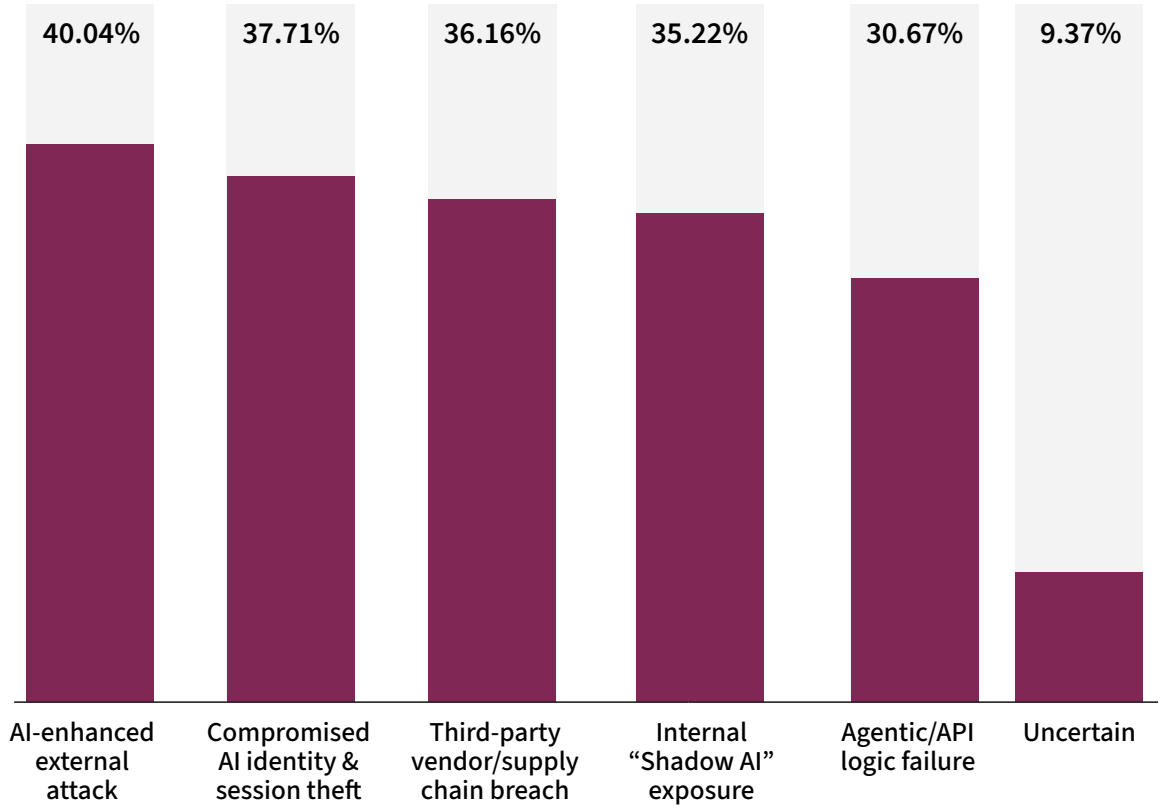

ASSESSING ENTERPRISE AI RISK: 5 CRUCIAL QUESTIONS FOR MODERN SECURITY TEAMS

1 Are our vulnerability scanners and threat intelligence feeds flagging rapid, automated exploitation attempts on newly released CVEs?

2 Where are our AI platform API keys, service accounts, and session tokens stored, and are we monitoring them for anomalous access?

3 What third-party AI plugins, browser extensions, or vendor integrations are active in our environment, and what permissions do they hold?

4 What are the top unapproved AI domains receiving traffic from our internal network, and which specific user groups are generating it?

5 Do we have comprehensive audit logging for internal AI agents, and can we detect when an agent executes an anomalous system or network change?

Dominant Adversaries Scaled AI into Operations

The research indicates a highly active threat landscape driven by specialized ransomware operations.

When asked which actors were detected in their networks over the last 12 months, LockBit (21%) and RansomHub (20%) claimed the top spots for the second year in a row.

Recent cybersecurity advisories indicate active threat groups like RansomHub are using AI to scale operations for:

- **AI-driven translation:** Instantly translating extortion demands into dozens of native languages.
- **Targeted phishing:** Crafting hyper-personalized threat letters directly to corporate board members.
- **Intelligent data parsing:** Using AI to autonomously scan stolen data, pinpointing high-value targets (like SSNs or medical records) to maximize extortion leverage.

APT 41, on the other hand, was only detected in 13% of networks, representing a near 50% decline in presence year-over-year.

APT 41's lower detection rate coincides with an operational strategy that remains notably less reliant on AI than many of its peers. APT 41 uses generative AI the way someone might use an administrative assistant, summarizing network topologies or translating technical data.

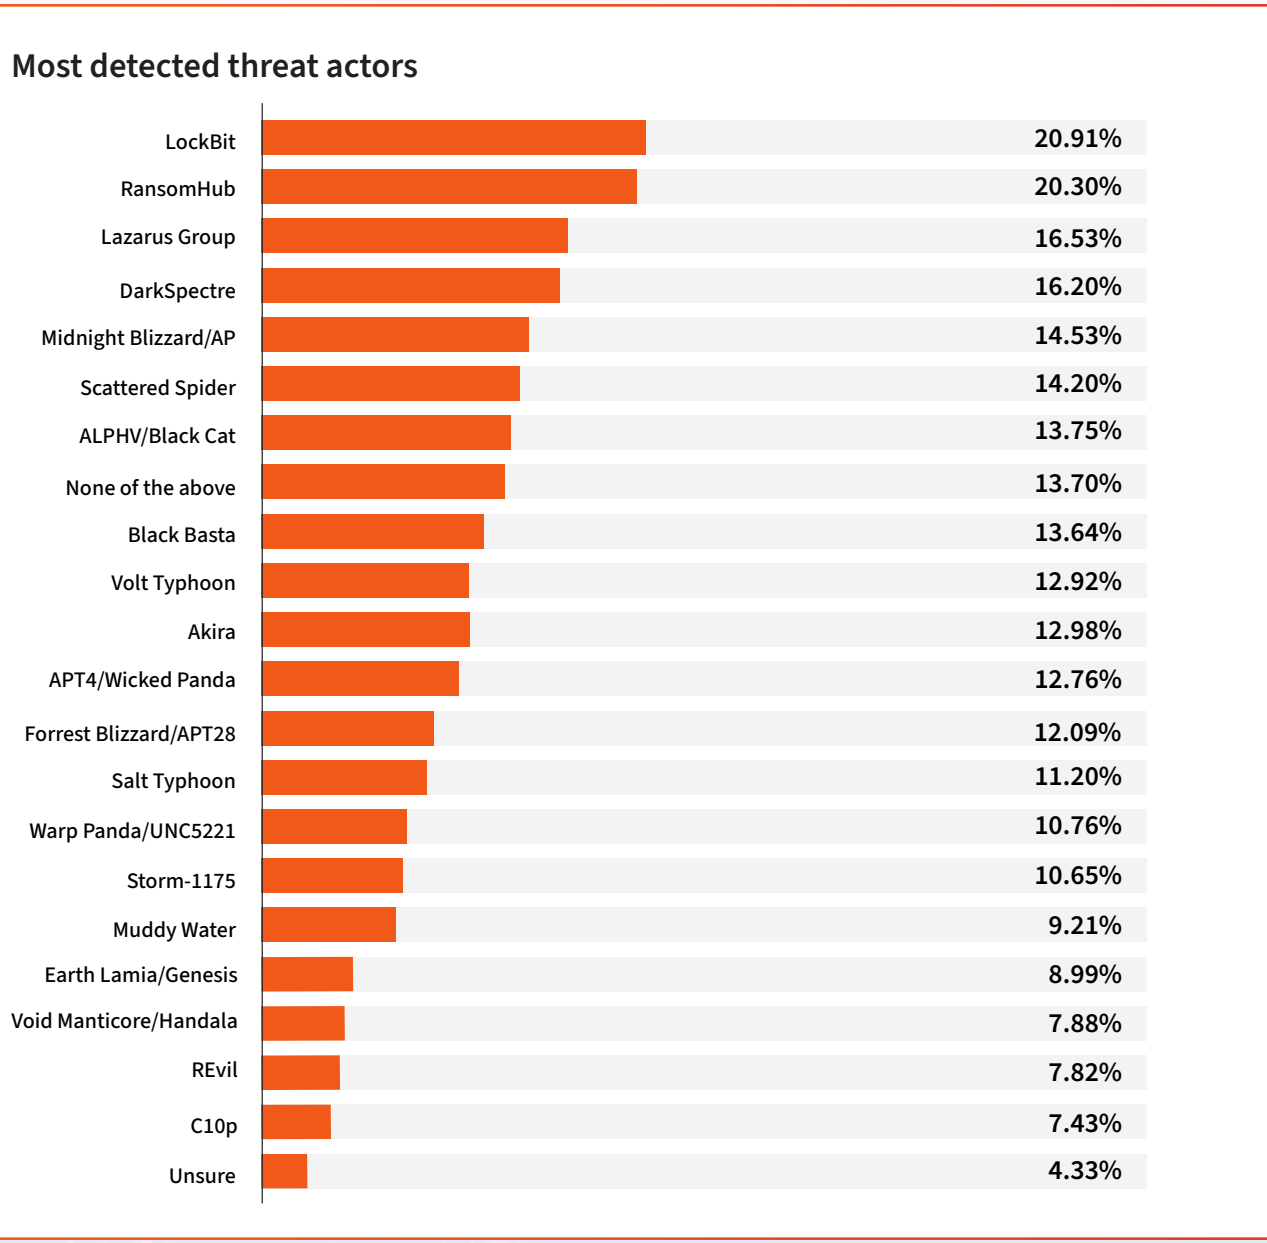

Post-Exploitation & Tactical Camouflage

Infiltration and Initial Access Trends

Initial access patterns remained consistent year-over-year, with no meaningful change in infiltration methods.

Phishing and other forms of social engineering (35.8%) remain the most common point of entry for attackers targeting organizations.

Software vulnerabilities follow as the second-most common initial point of entry (22.3%). Meanwhile, compromised credentials or brute-force attacks account for 14% of initial network infiltrations.

Most common initial point of entry for attackers

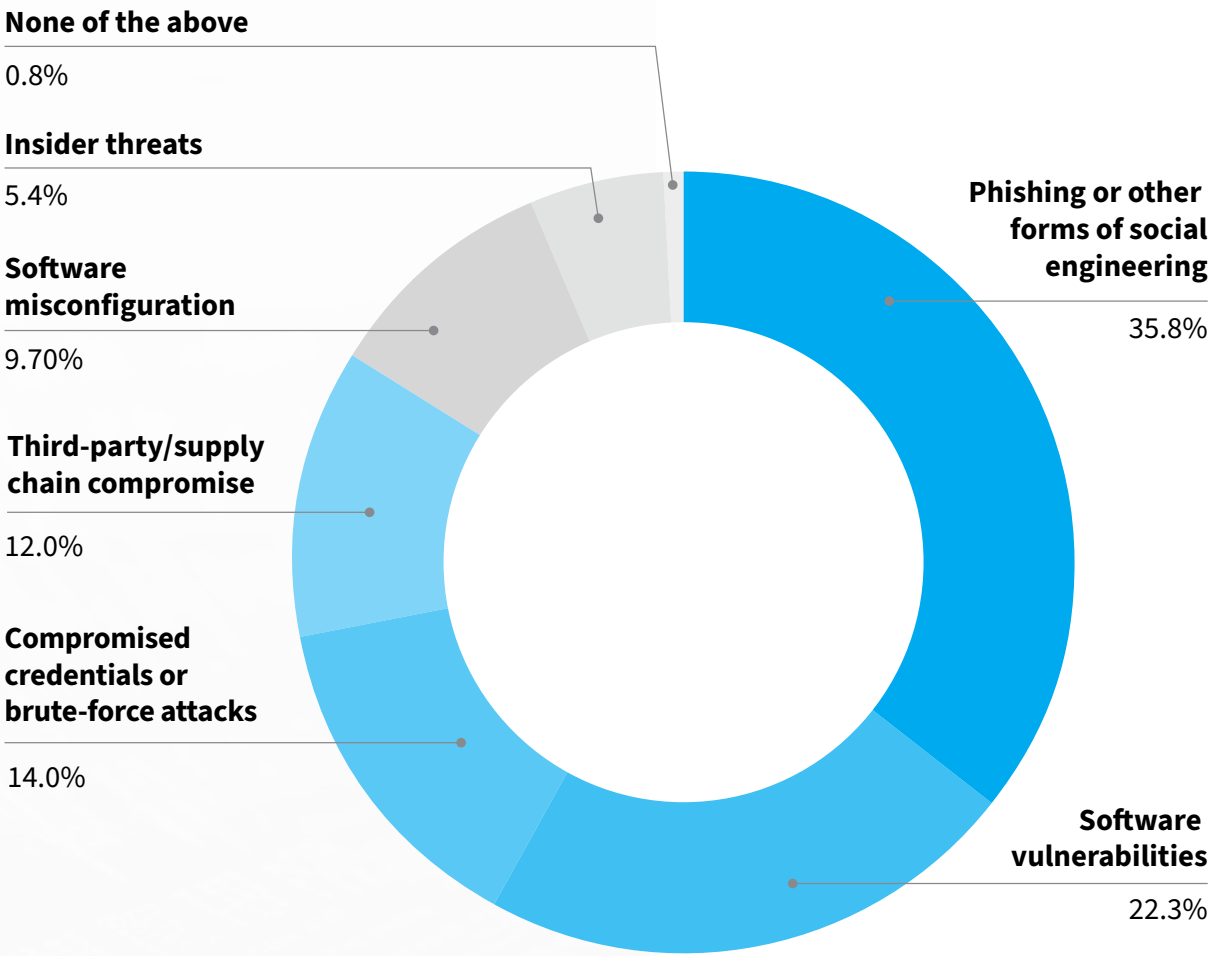

Internal Evasion and Obfuscation

Once inside the network, attackers are increasingly using stealth tactics to extend access, evade detection, and obfuscate malicious activity.

When asked how many distinct security incidents or 'near-misses' their organization had experienced involving compromised credentials, Living-off-the-Land (LotL), and encrypted evasion, respondents noted a growing increase in identity-based threats.

Compromised credentials	Living-off-the-Land (LotL)	Encrypted evasion
3.0	2.7	2.8
avg. incidents per org	avg. incidents per org	avg. incidents per org

COMPROMISED CREDENTIALS

Compromised credentials emerged as the most commonly observed evasion technique, averaging 3 incidents per organization globally.

The frequency of these incidents tracked highest in the technology sector and, separately, in the United States.

In September 2025, Scattered Lapsus$ Hunters used stolen credentials to infiltrate Jaguar Land Rover's networks, halting global production for five weeks at an estimated cost of £50 million weekly. The resulting bottleneck was so severe that the Bank of England cited the breach as a drag on UK GDP growth.

LIVING-OFF-THE-LAND

Living-off-the-Land (LotL) techniques followed closely behind compromised identities, recording a mean of 2.7 incidents per organization globally.

Living-off-the-Land (LotL):
A cyberattack tactic where adversaries weaponize an organization's own legitimate administrative tools to evade detection.

ENCRYPTED EVASION

Encrypted evasion tactics mapped to an identical pattern, also averaging 2.8 incidents per organization.

Looking Ahead

Artificial intelligence is shifting these evasive tactics from human speed to machine speed, drastically lowering the barrier to entry for highly sophisticated attacks.

Automated identity exploitation
Large language models (LLMs) will allow attackers to scale credential theft through localized phishing campaigns, while AI-driven sorting instantly identifies high-value accounts.

Context-aware LotL
Rather than guessing which tools to use, offensive AI agents can dynamically map an environment and generate custom native scripts that mimic the behavior of local sysadmins.

Polymorphic encrypted evasion
Attackers will leverage AI to mutate encryption protocols and traffic patterns on the fly, bypassing traditional network filters that rely on static signatures.

Detection Blind Spots

When asked which factors delayed a critical alert from being detected or investigated, respondents overwhelmingly pointed to the stealth tactics used by threat actors.

Organizations noted that in 41% of cases, attackers used encrypted channels to bypass detection. These channels utilize standard security protocols like HTTPS and TLS, which are designed to protect legitimate data privacy, but are hijacked by threat actors to conceal malicious communication and data theft.

Respondents also reported that 38% of attacker activity mirrored legitimate, authorized workflows and processes. This tactic involves weaponizing legitimate, pre-installed administrative tools and approved applications to execute tasks, ensuring malicious actions blend seamlessly into the background of an organization's daily business.

Another 34% also noted valid, high-privilege account permissions were used in the attack. This allows threat actors to bypass standard permission constraints entirely, utilizing legitimate administrative access to clear security logs, create backdoors, and reach high-value data targets completely unchallenged.

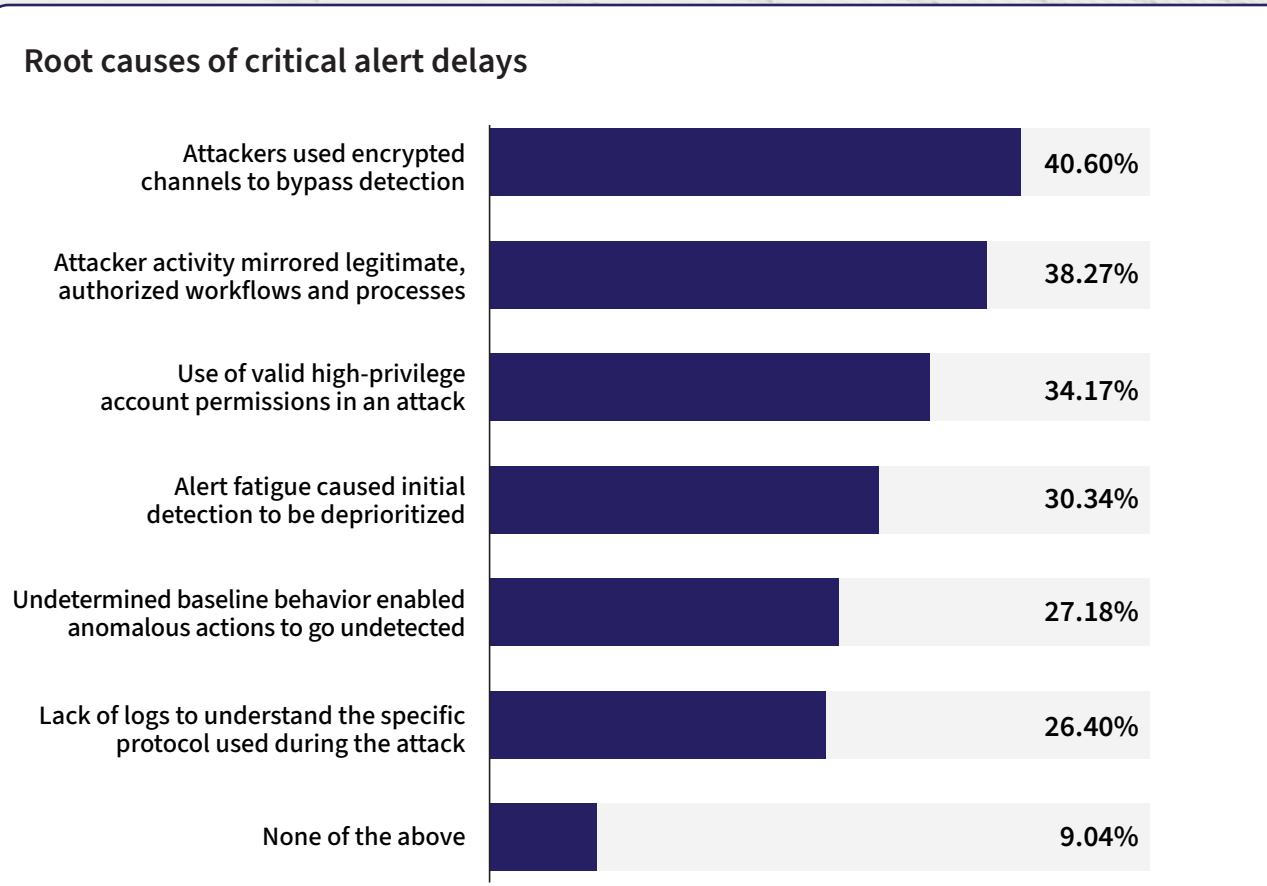

UNMASKING THE INVISIBLE THREAT

To stop adversaries from hiding in plain sight, organizations must counter each stealth tactic with a specific, targeted defense.

Encrypted evasion

Deploy inline SSL/TLS decryption to inspect traffic in transit. Unmasking the actual data within encrypted streams is critical to exposing hidden malicious activity before it can execute.

Workflow camouflage

Continuously baseline the normal operational rhythms, timing, and data volumes of your environment. By understanding what “normal” network behavior looks like, organizations can instantly flag the subtle, anomalous deviations that occur when a legitimate tool is weaponized.

Privilege abuse

Monitor privilege use directly on the network in real time. By tracking how and where elevated credentials move across the wire, security teams can immediately detect abnormal resource access and lateral movement to catch the threat the moment it touches the network.

The Ransomware Economy & The Dwell Time Crisis

Ransomware Volume and Geographic Shifts

Surveyed organizations reported an average of 3.5 ransomware incidents over the last 12 months, a decline from last year's 5.4 incidents per organization.

On a regional level, Germany reported the highest average number of incidents in the survey, with 4.8 attacks per organization.

Across all surveyed regions, agriculture (4.5) and telecom (4.2) emerged as the most frequently targeted industry sectors, experiencing the highest global averages.

THE RANSOMWARE MIGRATION

Ransomware isn't shrinking, it's migrating.

As coordinated global law enforcement hardens traditional targets, syndicates are moving downstream to other targets.

	Primary Target Markets	Emerging Growth Markets
Regions	United States, Western Europe, Middle East & Africa, Australia, Singapore	Latin America (Brazil, Mexico) & APAC (Thailand, Indonesia, Vietnam)
Operational Status	Hardening via aggressive law enforcement interventions (such as the FBI, NCA, Europol, and ASD)	Vulnerable due to rapid enterprise digitization outpacing regional cybersecurity infrastructure maturity

Post-Compromise Detection Delays

While overall attack volumes appear to be declining, the data reveals a troubling counter-trend: the timeline to identify threats has lengthened.

Organizations are now documenting a distinct delay between initial access and actual ransomware detection compared to previous years.

When asked at what point they recognized their organization was being targeted by a ransomware attack, nearly half (49.3%) of organizations only identified ransomware activity after it reached the data exfiltration or later stages (encryption, ransom demand).

In 2025, that figure was just under a third of organizations.

Nearly 15% did not recognize they were being targeted until a ransom demand was issued. This represents a dramatic increase from the previous year, which hovered around 6%.

Attack stage at point of ransomware recognition

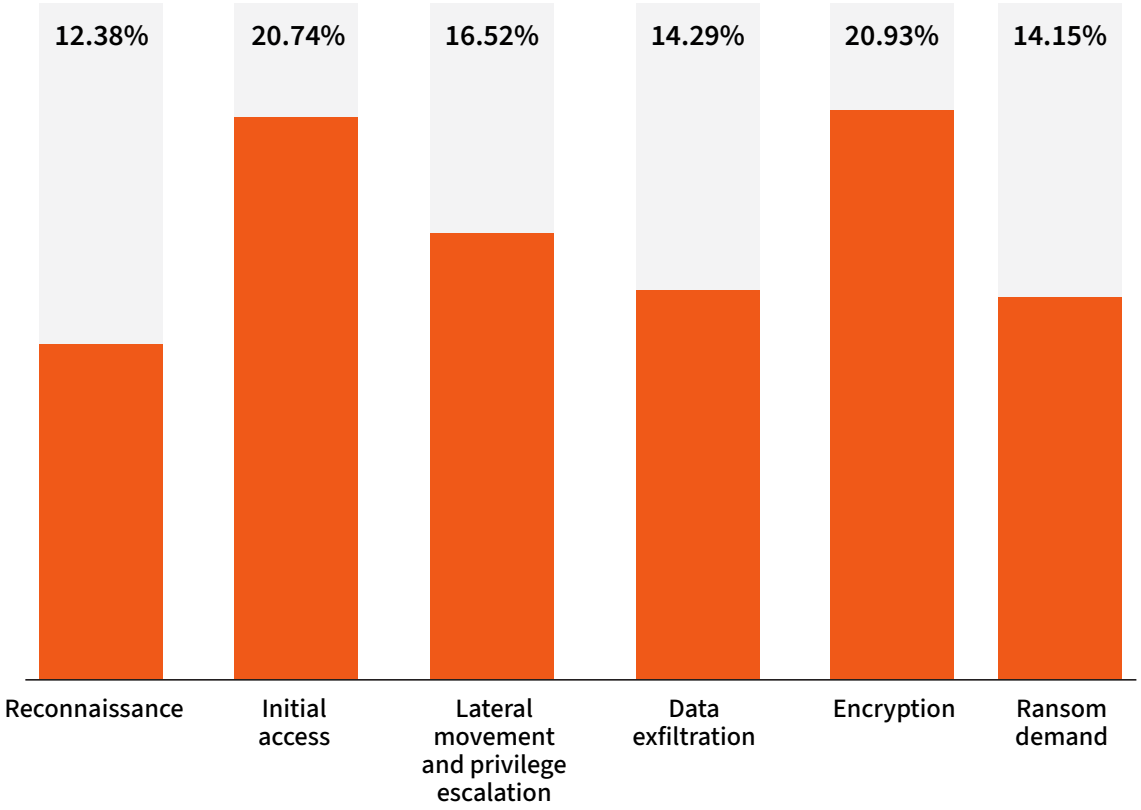

Extended Attacker Dwell Time

When asked how long a threat actor maintained access to systems following a ransomware incident, respondents reported a mean dwell time of 2.4 weeks.

The duration reflects a year-over-year increase from the prior year's report, where organizations cited an average dwell time of 2 weeks.

Every hour ransomware goes undetected drastically increases its potential blast radius. A wider detection window grants adversaries the critical dwell time needed to move laterally and locate backups. This delays containment, turning what could have been a localized incident into an organization-wide crisis.

Threat actor dwell time following ransomware incidents

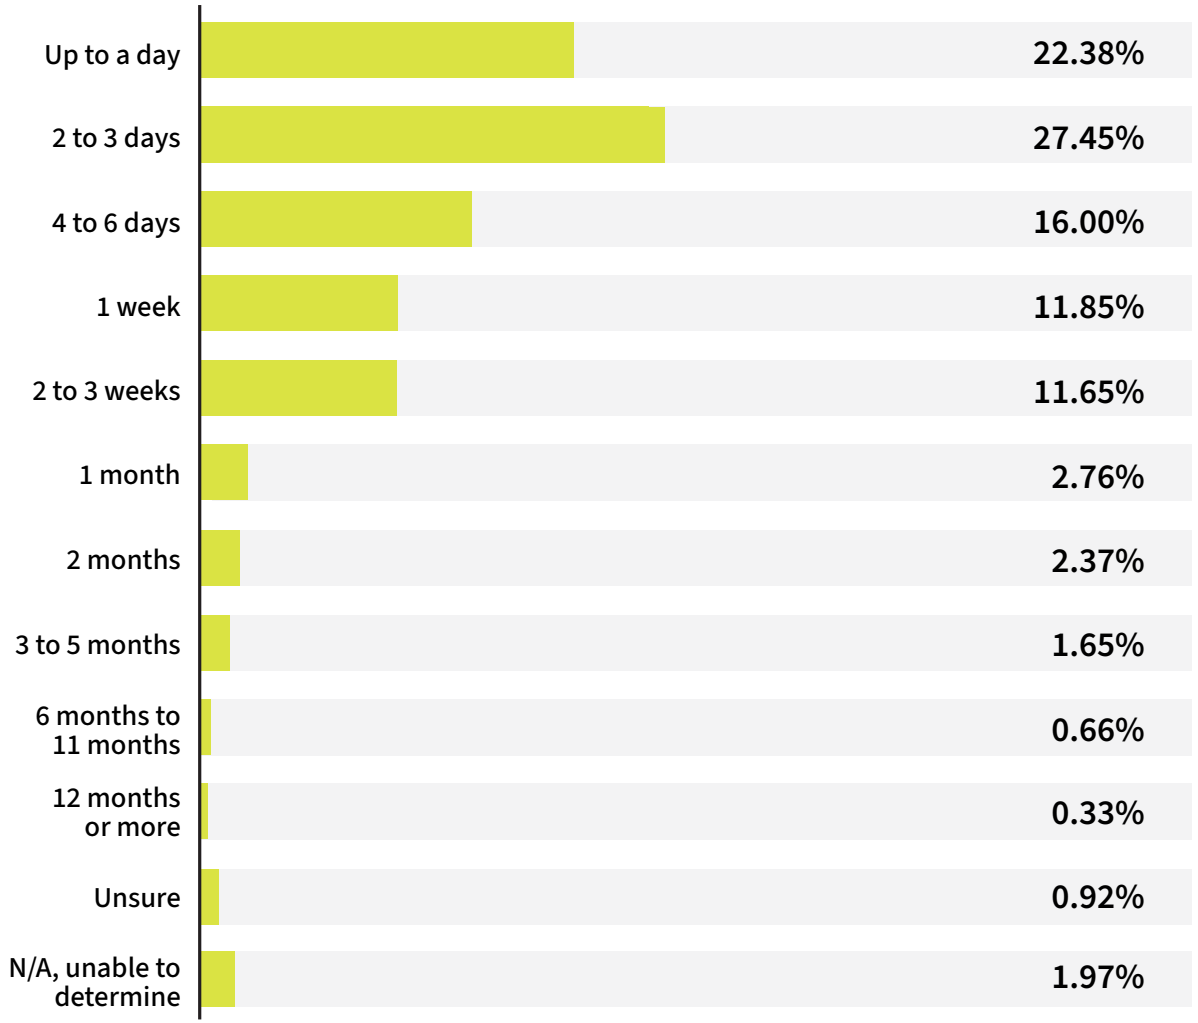

Evolving Extortion Economics

More organizations are paying ransoms than before.

In this year's survey, 83.5% of respondents reported that their organization paid a ransom, up from 70% in the prior year.

On average, the total cost of ransomware payments amounted to $2.8 million this year, a decline from $3.6 million in 2025.

Top payouts were seen in France at $4 million, followed by Germany at $3.7 million and the UAE at $3.4 million.

Total ransomware payments by country

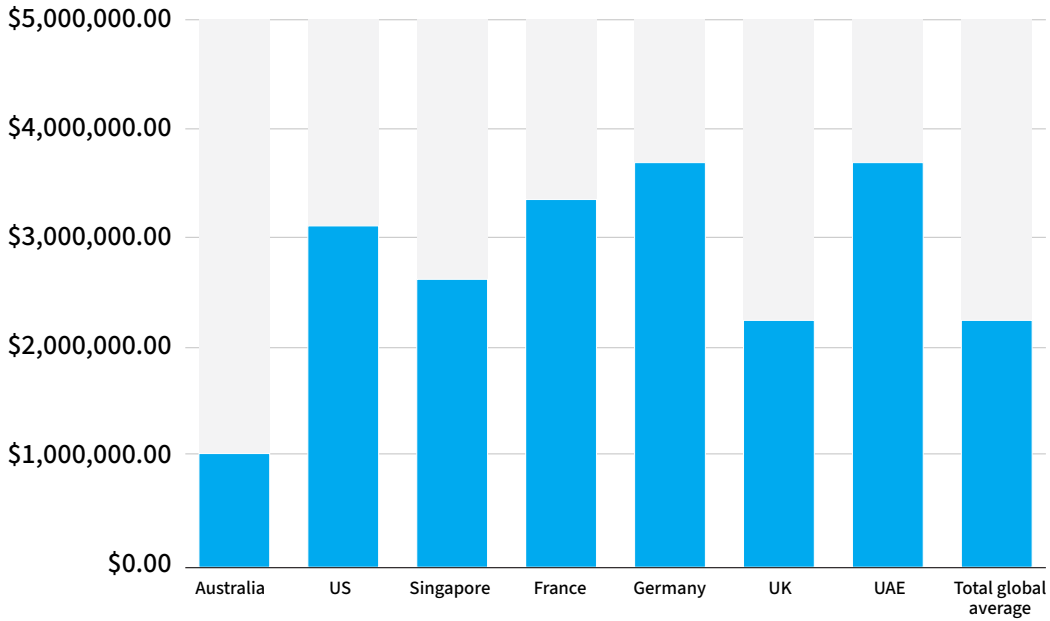

Threat Intelligence Note:
Extended attacker access often gives adversaries more time to identify, exfiltrate, and encrypt high-value data. In some cases, this prolonged operational window leaves organizations with no viable alternative to payment.

BEYOND THE DATA: WHAT'S DRIVING THESE TRENDS?

While the survey highlights shifting ransomware outcomes, it doesn't explicitly capture the why. Based on broader industry indicators, we attribute these shifts to two likely factors:

- **Accessible price points:** Lower ransom demands may paradoxically increase payment likelihood by broadening the pool of organizations capable of paying.
- **The professionalization of response:** A maturing ecosystem – including cyber insurance, incident response firms, and professional negotiators – is enforcing a more structured, capped approach to settlements, ultimately driving down per-incident payouts.[1]

1. https://web.coalitioninc.com/download-2025-cyber-claims-report.html

Operational Disruptions

Beyond the immediate financial impact, security incidents continue to cause operational disruptions.

When asked how much downtime their organization experienced following a cyber incident, respondents reported a global average of 29.5 hours, with peak rates occurring in the U.S. and Australia.

This marks a decline from the 37.2 hours of average downtime documented in 2025.

THE PRODUCTIVITY PREMIUM

Shorter outages do not automatically equal lesser impacts. Because the increased use of AI in the enterprise translates to higher baseline productivity, any period of disruption likely deals a heavier material blow to total output than in the past.

Average downtime per incident

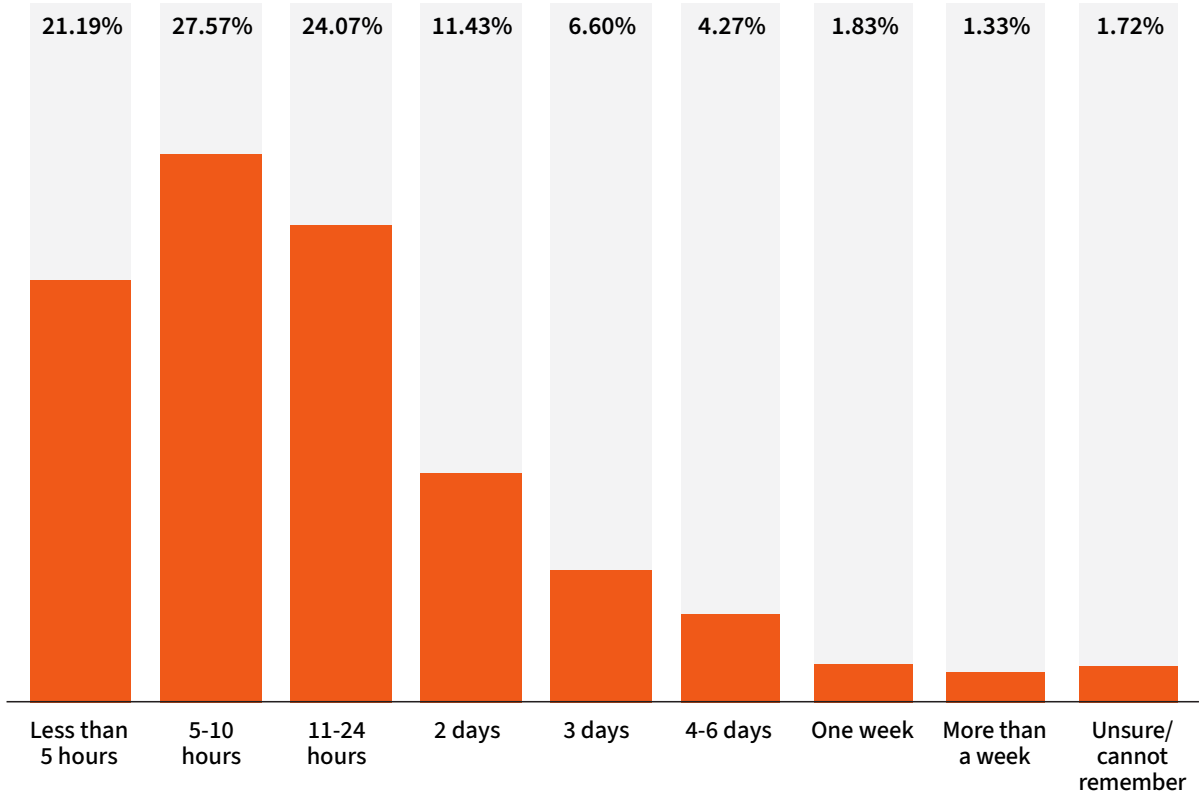

Friction on the Path to the Agentic SOC

The Automation Gap

SOC workflows remain heavily manual, with most respondents reporting mid-to-high levels of manual intervention across the threat lifecycle.

When asked what level of manual intervention is required across different stages of the threat lifecycle (threat detection, alert triage, investigation, and response), respondents reported an average of 45%.

Investigation was the most manual stage, requiring manual intervention about 50% of the time.

Mean manual effort required across the SOC threat lifecycle

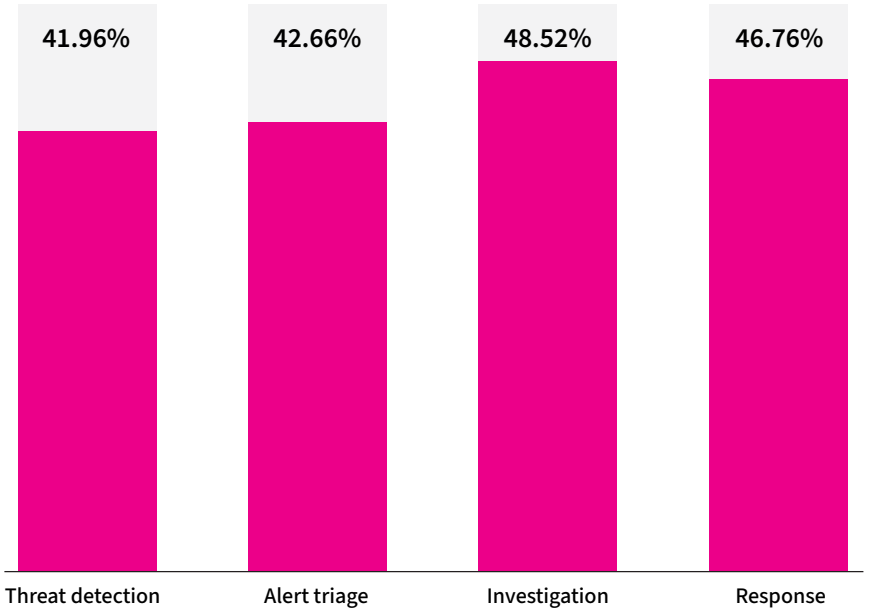

THE INVESTIGATION PARADOX IN THE AGE OF AI

Even as security technologies advance, the investigation phase remains a heavily manual bottleneck. Traditional alerts typically lack deterministic context, leaving analysts to manually correlate user behavior, privilege levels, and encrypted data streams to validate a threat. In an era where AI can accelerate operations, automation can only succeed if it is fed rich, comprehensive context that turns disjointed alerts into instantly readable, end-to-end attack timelines.

Proactive Security Constraints

Bogged down by manual data gathering and triage constraints, SOC analysts are limited to spending just 44% of their time on proactive efforts, like threat hunting and detection engineering.

An imbalance can force teams into a reactive posture, creating a dangerous operational gap where sophisticated, slow-moving threats can dwell undetected inside the network.

The Containment and Resolution Timelines

When tracking the full lifecycle of a security incident, respondents reported that alerts take an average of about 2 weeks to move from initial detection to full resolution or containment.

This represents a slight improvement from 2.3 weeks from prior year. However, the timeline stretches even further for the U.S., where containment averages 2.4 weeks.

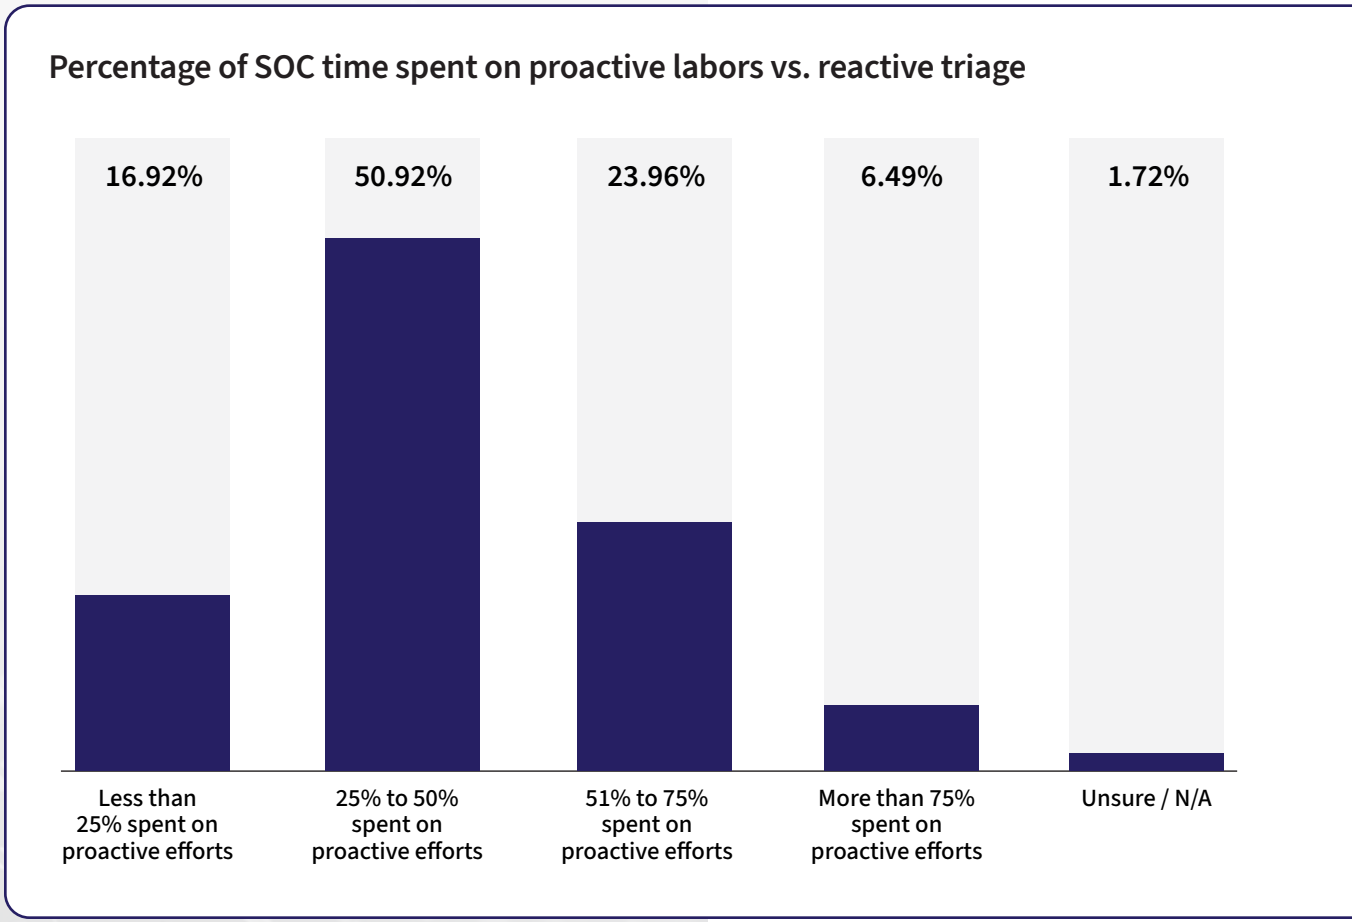

The AI Friction Point

As manual workloads and response delays persist across SOC operations, organizations are expanding investments in agentic capabilities to improve detection and response efficiency.

However, AI-enabled security tooling is also introducing measurable investigation delays, with false positives impacting investigation timelines across the majority of SOCs.

When evaluating the specific operational impact of these alerts, respondents reported that AI-generated alerts or detections led to false positives that negatively impacted investigation timelines nearly 30% of the time, on average.

THE AUTOMATION TRAP: FASTER RESPONSE TO THE WRONG THREATS

Automating triage before fixing data quality doesn't solve the noise problem… it just accelerates it.

- **The illusion of speed:** Shaving minutes off your response time means very little if autonomous agents are just chasing ghosts at lightning speed.
- **The ground truth:** High-fidelity network context is the only reliable baseline for trusted automation, separating sophisticated attacker behavior from legitimate user activity.
- **The bottom line:** Scalable SOC automation lives or dies by signal integrity. If the underlying detections can't be trusted, faster execution isn't an efficiency gain, it's a liability.

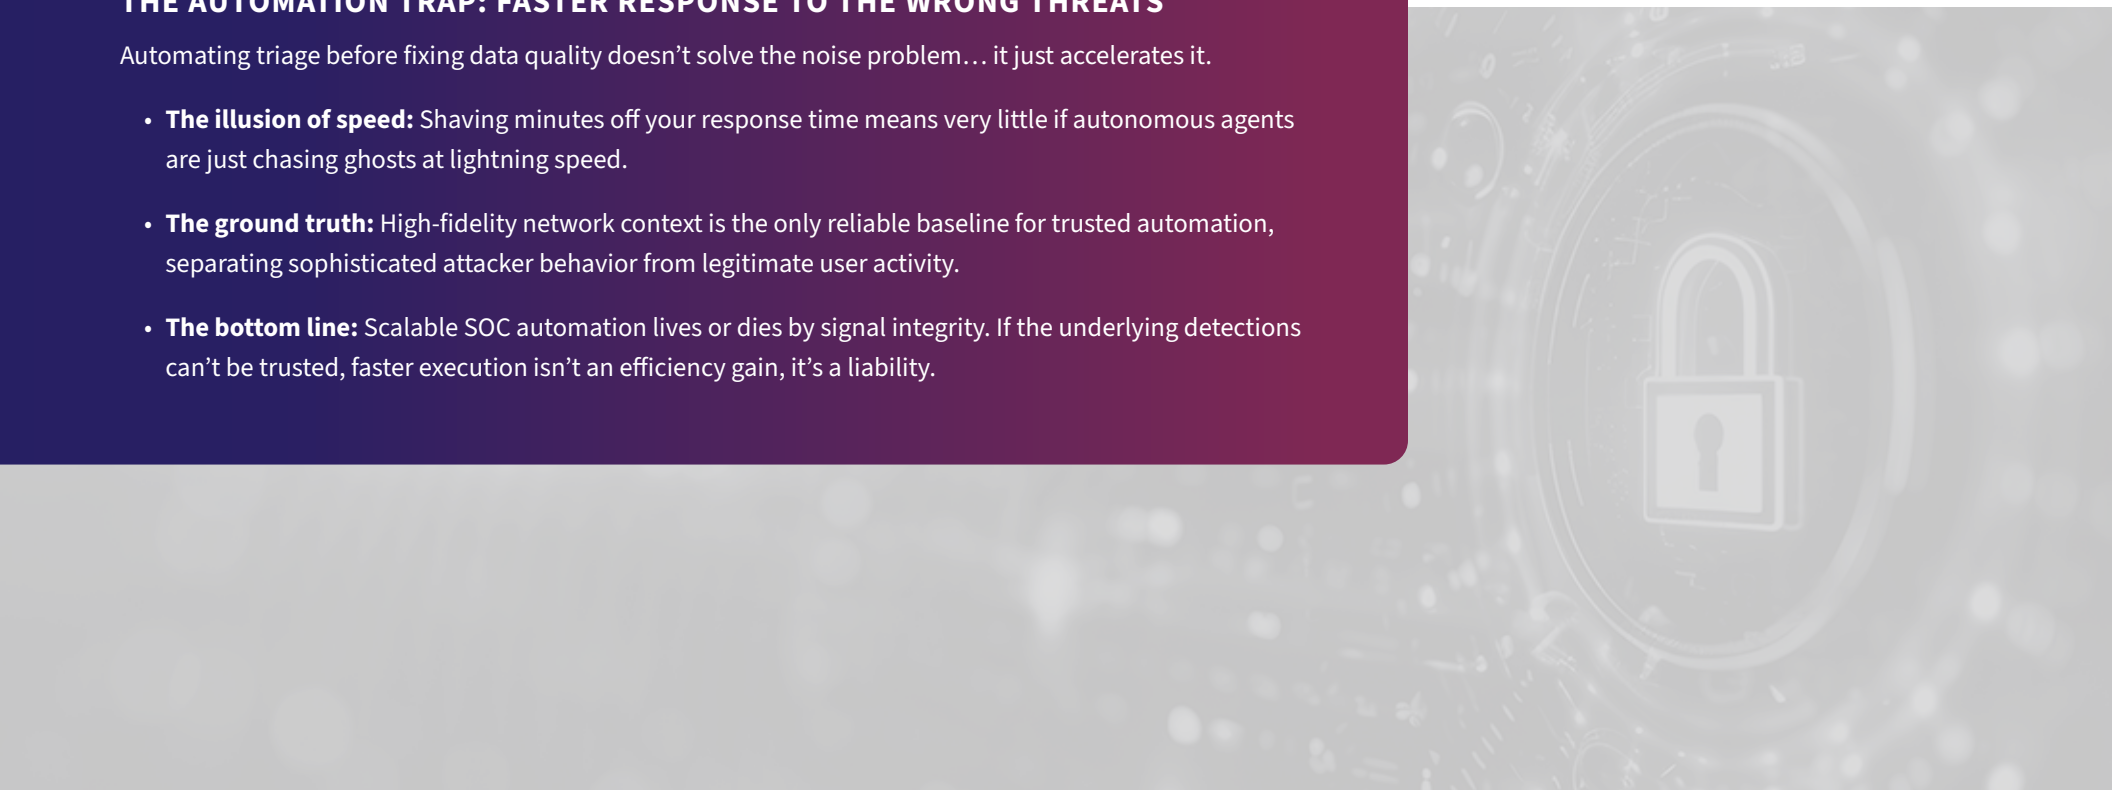

Defending the Enterprise in the Age of AI
3 Steps to Take Today

1. REDEFINE YOUR ATTACK SURFACE

Inventory and map your expanding AI ecosystem – generative apps, autonomous agents, and third-party integrations – to secure the modern, non-human entry points attackers are weaponizing.

2. OUTSMART EVASIVE THREATS

Shift from static signatures to continuous behavioral baselines. To expose AI-mutated scripts and stealthy Living-off-the-Land tactics early in the cycle, you must detect the subtle anomalies that perfectly mimic legitimate traffic.

3. FUTURE-PROOF YOUR SOC

Transition to an agentic SOC fueled by high-fidelity network context. By automating detection, investigation, and response to close the velocity gap, you elevate human analysts from manual responders to strategic AI guardians.

Methodology

In conjunction with Censuswide, ExtraHop surveyed 1,803 security and IT decision-makers (director level or above) working for organizations with 1,000+ employees in the U.S., U.K., France, Germany, Singapore, Australia, and the United Arab Emirates in May 2026.

About ExtraHop

ExtraHop turns the network — the enterprise's ultimate source of truth — into actionable insight to power security, performance, and resilience. Delivering superior data by design, we ensure superior defense by default.

The ExtraHop modern network detection and response (NDR) platform provides visibility that thinks, analyzing behavior to intercept evasive risks before they cause damage. We transform network noise into definitive context, enabling security teams to make faster decisions and operate at uncompromised scale.

Whether securing cloud modernization or de-risking AI adoption, ExtraHop gives global enterprises the ground truth they need to thrive.

To learn more, visit extrahop.com or follow us on LinkedIn.

EXTRAHOP®