Global Security Report

The AI Speed Tax

Why the fastest-moving businesses are slowest to recover

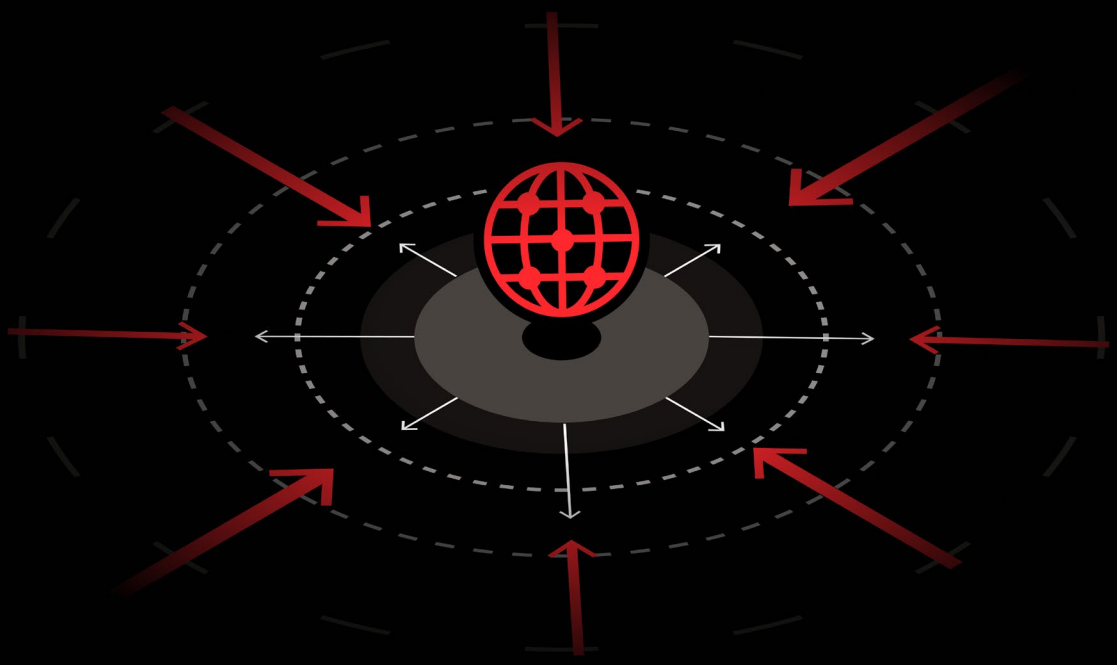

fastly

Table of Contents

01 Executive Summary

03 The AI-first Paradox

03 AI creates its own attack surface
04 How AI scrapers hit the bottom line
04 WAAP becomes essential

05 When bugs beat breaches

05 Software bugs actually triggered

06 A silver lining in incident recovery

06 Investment in response capabilities pays off

07 Investment priorities: where security dollars are really going

07 With mounting risks, it's no wonder that data protection and privacy lead investment priorities at 37%

08 CISOs have more responsibility but hollow support

08 Policy changes miss the point
08 Nobody knows who's in charge

09 Solving the skills strain

09 Alternatives to external recruitment
10 How threats vary by sector

11 The path forward: security by design in an AI-accelerated world

11 The automation ceiling
11 Two problems, one solution
12 How to fix security before you start

13 About the research

Executive Summary

Businesses are hurtling towards a cybersecurity crisis of their own making because of the massive scale of adoption of AI within organizations. By embracing AI innovation without thinking through how to reinforce their security and putting a strategic plan in place to implement holistic solutions, companies racing to call themselves "AI-first" are discovering that shiny new functionality without infrastructure defense creates more problems than it solves.

To understand what companies are experiencing in the real world, beginning in September 2025 Fastly partnered with research agency Sapio to survey 2,000 IT decision makers across 21 regions who are involved with cybersecurity. The findings reveal an uncomfortable truth: the organizations most aggressive about AI adoption are the ones struggling the most with security incidents.

Here's what emerged:

Businesses are paying a steep AI tax.

More than 75% of businesses that identify as AI-first (meaning they've integrated AI into core processes from the outset, either publicly or informally) take an average of 80 days longer to recover from security incidents than their peers. It takes 6.8 months on average across all regions for AI-first organizations to recover, and 3.9 months for everyone else.

The financial impact is also more painful for AI businesses: incidents cost AI-first organizations 135% more than their counterparts, consuming 3.13% of annual revenue compared to 1.33% for non-AI-first companies -representing billions of lost revenue. Almost half (44%) report that AI was directly exploited in their most recent incident.

continued on next page

The expanding attack surface: AI is expanding faster than security can follow

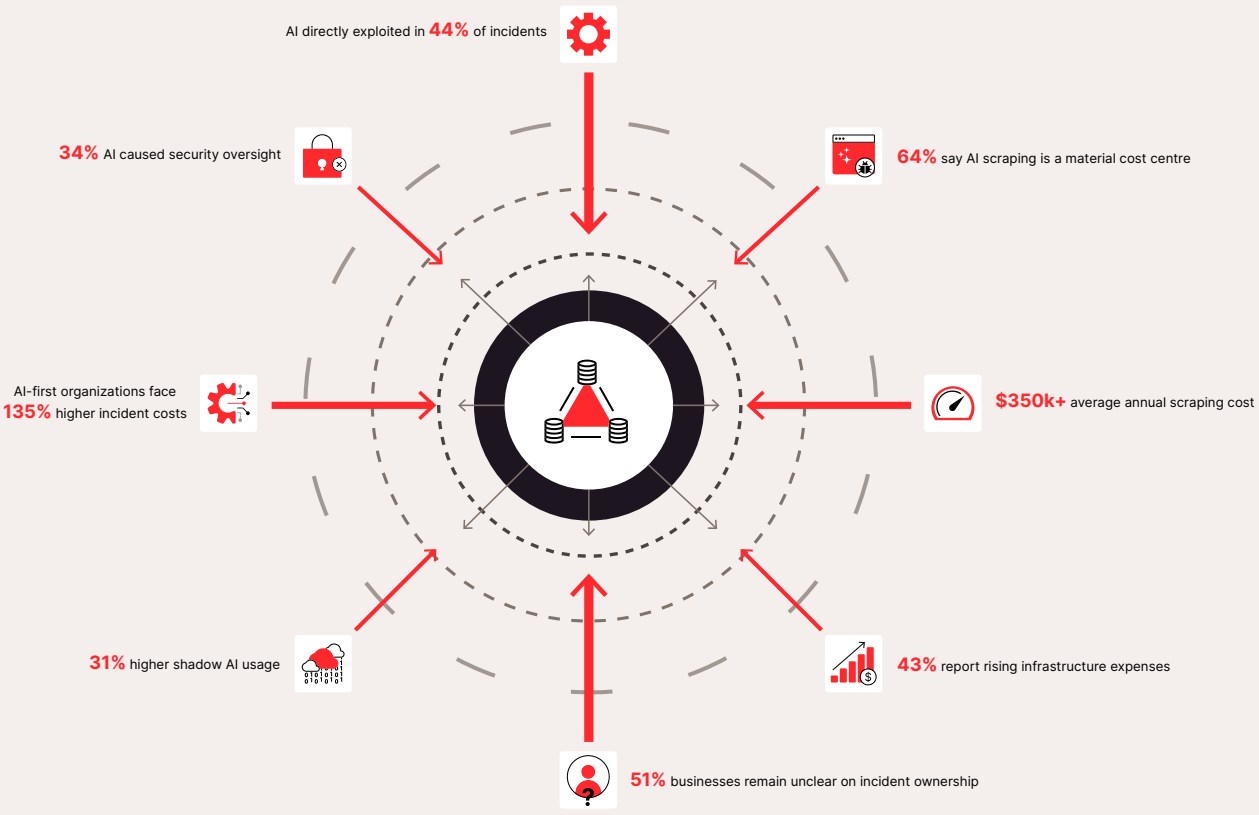

The infrastructure footprint is expanding faster than defenses can keep up. Shadow AI also jumped in the past 12 months. More than a third (34%) of AI-first organizations say direct exploitation of AI contributed to their last incident, with another 30% saying that AI use led to an oversight that contributed to the incident. Meanwhile, AI scraping has become a material cost center for more than two-thirds (64%) of businesses, with average annual infrastructure costs rising by almost $350,000. Further, 43% of respondents reported surging infrastructure expenses. It isn't just a financial issue: 40% also faced operational disruption and 29% reported degraded customer experiences.

Nobody knows who's actually responsible when things go wrong. Half of AI-first businesses (51%) report confusion over who handles incident response, compared to 23% of non-AI-first organizations. Traditional accountability is rapidly deteriorating because teams can now include humans and self-thinking, self-learning, and self acting AI agents.

There is a bright spot in the data. Overall incident volume for all businesses held steady at an average of 41 incidents per organization, and recovery capabilities improved. Average recovery time dropped more than a month from last year's 7.34 months to 6.08 months. Revenue losses from incidents fell to 2.68% of annual income, down from 2.98%. Organizations that invested in post-incident reviews (52%) and response automation (43%) are likely among those seeing positive results.

Investment priorities are shifting to address AI-specific risks. Agentic discoverability tools lead security spending on agentic AI infrastructure at 56%, followed by API security (55%) and web application firewalls (54%). Three-quarters (75%) worry about DDoS attacks targeting AI agents, while 53% acknowledge they lack AI-specific security expertise.

The path forward requires a mindset shift. Enterprises must monitor AI crawler activity, anticipate shadow AI adoption, and strengthen perimeter defenses before expanding the attack surface. It's also prompting a rethink of web application and API protection (WAAP) solutions; companies are starting to view them as business-critical infrastructure rather than niche solutions.

AI-first companies tend to suffer from AI-related compromises

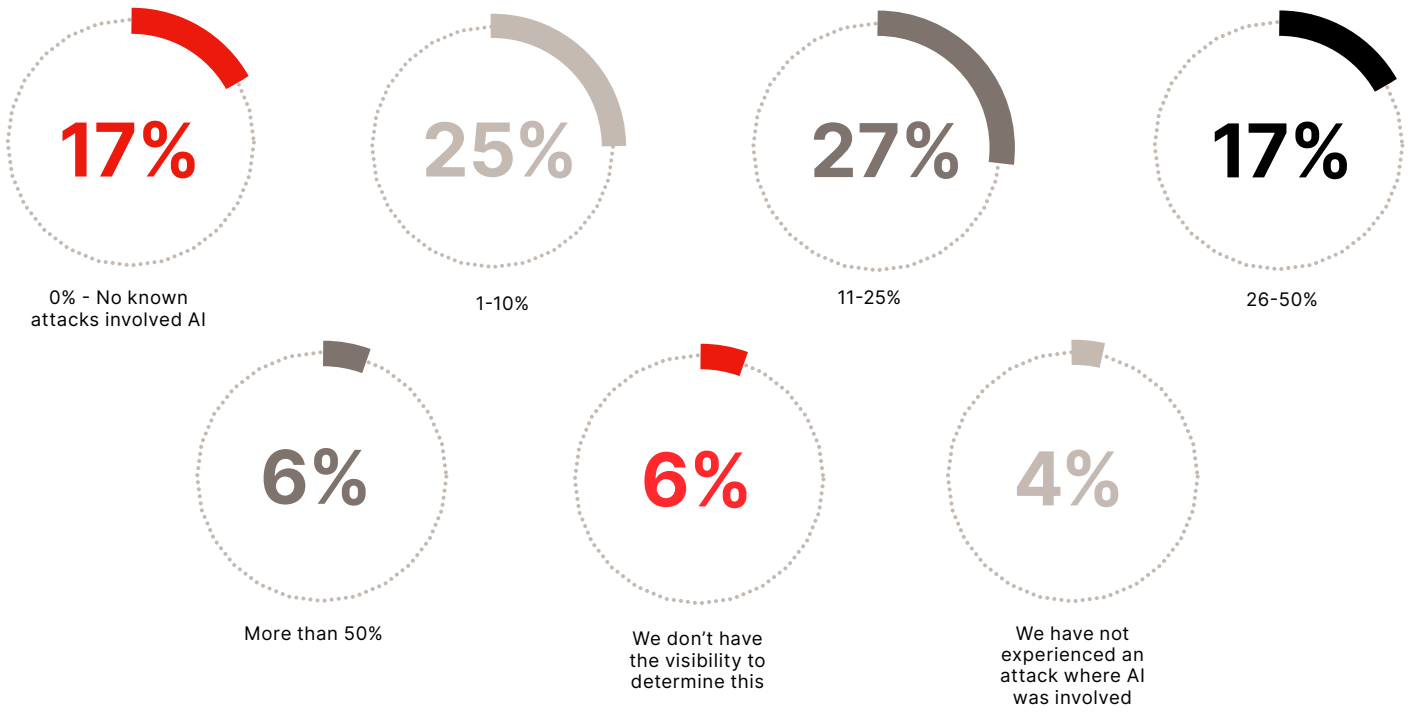

The AI-first paradox

AI is a potential productivity booster, but as with any enabling technology, it requires a strategic balance between innovation and security. Organizations that brand themselves as "AI-first" and race to integrate artificial intelligence must find this balance to truly apply AI's benefits. At the moment, these businesses take an average of 6.8 months to recover from cybersecurity incidents, which is 80 days longer than their non-AI-first peers, who recover in just about 4 months. They have some work to do.

> "The AI tools themselves are going to be privileged parts of your infrastructure, and that's what's going to create the risk."
>
> — Marshall Erwin, CISO at Fastly

Innovation is a critical business advantage, but organizations that create robust security measures to protect AI data, infrastructure, and processes will be better positioned to innovate than those who don't. Incidents cost AI-first businesses 135% more than non-AI-first organizations. The recovery gap translates directly into lost revenue, extended downtime, and prolonged reputational damage.

AI creates its own attack surface

AI introduces more complexity, more code, and more infrastructure that companies must protect. Almost half (44%) of AI-first organizations report that AI was directly exploited in their most recent security incident, compared to a mere 6% among non-AI-first businesses.

One apparent cause traces back to what security teams can't see. Shadow AI (unauthorized tools that employees adopt without IT approval) runs 31% higher among a quarter of employees at AI-first organizations, presumably because the culture encourages innovation with AI. But for Marshall Erwin, CISO at Fastly, sanctioned AI tools are at least as much of an issue as shadow AI.

Approved AI tools often receive extensive automated permissions, and companies already struggling with identity and access management in a pre-AI world are watching that problem explode. "These tools are going to expand potential access risks," Erwin warns. "The AI tools themselves are going to be privileged parts of your infrastructure, and that's what's going to create the risk."

The numbers support this conclusion. Over a third (34%) of AI-first organizations cite AI usage as a factor in security oversights that contributed to their last breach. That compares to 20% in traditional organizations.

Think of AI tools as machine entities requiring their own identity governance. Automated privileges allow automated attacks. The more agentic these tools get (meaning the more complex their autonomous tasks) the bigger that risk will become.

AI-first organizations took longer to recover from security incidents than non-AI-first ones

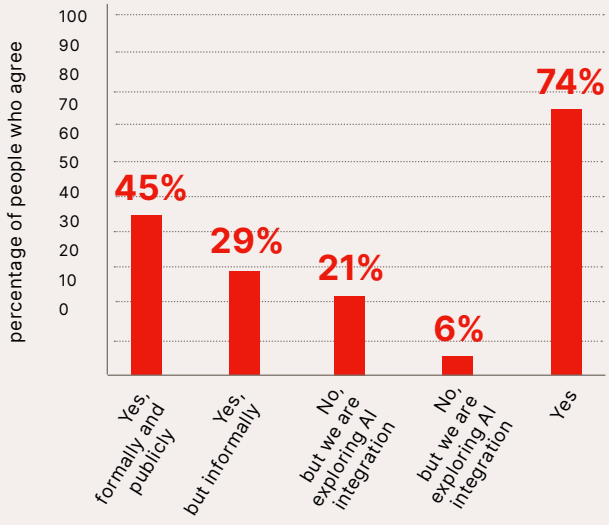

1 (employees using unsanctioned AI tools)

How AI scrapers hit the bottom line

Shadow AI invites accidental misuse of AI inside an organization, but there are also risks from third parties using AI. These malicious actors can target an organization's content using scraper bots. AI scraping is costing companies serious cash as it puts their infrastructure under strain.

Approximately two thirds of businesses (64%) say AI scraping has become a material cost center, with expenses soaring over $348,000 annually on average. That's hitting infrastructure bills and operational budgets hard. More than four in ten companies have watched infrastructure expenses climb as AI activity ramps up. Another 40% report operational disruption, while 29% are dealing with user experience problems. Sluggish load times, broken functionality, and degraded performance are the kinds of problems that send customers elsewhere if they persist.

Old-school external attack techniques can do serious damage to AI infrastructure. Three quarters of respondents to Fastly's survey worry about DDoS attacks hitting AI agents. What this means is that even companies who don't use AI internally should be mindful of how others might use it to exploit them.

None of these risks should stop companies from embracing AI, but the winners will be those who innovate while also adjusting their security postures accordingly. That means either finding the skills to do so (53% admit their security teams lack the AI specific expertise to deal with these threats) or working with a third-party partner to help manage the risks.

WAAP becomes essential

These realities are changing how organizations think about their security stacks. WAAP might once have been lower on the list of security tools for some companies, but it's now becoming part of their core infrastructure. It's the control layer for managing costs and securing the APIs that underpin modern digital services from those both inside and outside a company.

Enterprises are voting with their wallets. When they invest in protecting agentic infrastructure, organizations are prioritizing agentic discoverability (56%), API security (55%), and web application firewalls (WAFs, at 54%). These non-traditional security categories are direct responses to architectural patterns that barely existed two years ago.

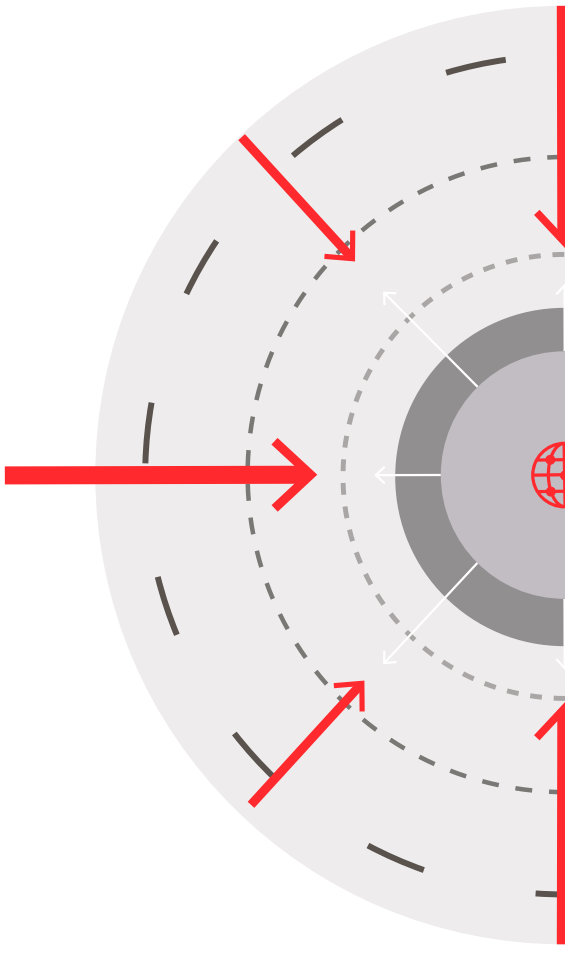

When bugs beat breaches

Our research showed that in 2025, the number of security incidents stayed flat. Organizations faced an average of 41 known incidents, up just one from 2024. But the headline number tells you almost nothing. What's actually breaking tells a more realistic story.

Software bugs actually triggered 40% of incidents, up from 33% in 2024, moving it from second place to the top slot, knocking the incumbent top cause (external attackers) down to second place at 39%. Misconfigurations took third place, up to 29% from fourth place at 25% last year. More than ever, companies are wrestling with security failures that have nothing to do with sophisticated threat actors and everything to do with how they write their code or configure their infrastructure (which is increasingly also done with code).

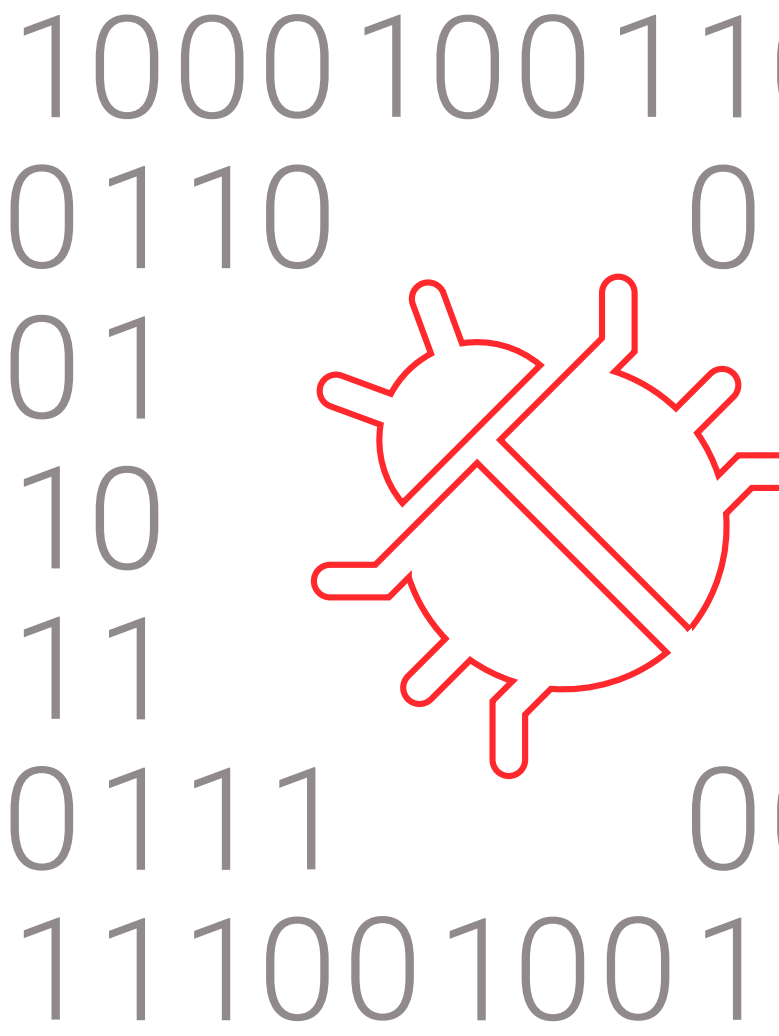

Software bugs actually triggered 40% of incidents, up from 33% in 2024.

Ninety percent of organizations suffered at least one cybersecurity incident, but it's clear that these stem from development problems, process problems, and maybe cultural problems in the enterprise. Clearly organizations cannot take their eyes off the perimeter, but it's time for them to elevate their focus on budget allocation and team structure rather than organizing purely around external threats.

Scale amplifies all these issues. Large enterprises with 10,000+ employees averaged 57 incidents, nearly 40% above the mean of 40. Sprawling attack surfaces and tangled development pipelines create more opportunities for things to break. Smaller organizations deal with the same issues at a smaller scale, which helps but doesn't eliminate the problem.

A silver lining in incident recovery

The Fastly survey revealed some genuinely good news: organizations are getting better at bouncing back from attacks. Average recovery time dropped to 6.08 months in 2025 from 7.34 months the previous year. That's more than a month off the recovery timeline, representing a meaningful improvement when every day of downtime costs money and erodes customer trust.

The gap between expectations and reality is closing too. Organizations now expect recovery to take 5.89 months, and the actual timeline of 6.08 months is remarkably close. This alignment suggests that companies are developing more realistic incident response plans based on actual experience rather than wishful thinking.

Financial impacts are also trending in the right direction. Revenue losses averaged 2.68% of annual income, down from 2.98% last year. While that's still painful (a mid-sized company losing nearly 3% of revenue is taking a serious hit), the trajectory matters. Organizations are containing the damage more effectively than they did twelve months ago.

Lastly, customers seem to be more forgiving once they see a company making efforts to fix its cybersecurity problems. Reputational recovery averages 4.73 months, faster than the 6.08 months needed to fully restore systems and operations. Managing the crisis narrative and maintaining customer communication can rebuild trust even while technical teams are still cleaning up the mess.

Investment in response capabilities pays off

Organizations are putting money where it matters. Over half (52%) invested in post-incident reviews, systematically analyzing what went wrong and how to prevent it next time. Another 43% have implemented response automation, using technology to speed up containment and recovery steps that previously required manual intervention.

But there are significant caveats buried in these positive trends. Despite the improvements, 30% of organizations still lack regularly tested incident response playbooks. The pressure is on to squeeze real recovery times further.

And while companies might recover quickly, lighting frequently strikes twice or even more. Two thirds (66%) of organizations suffered repeat incidents within three months, as underlying problems might still exist.

Recovery expectations are getting more realistic

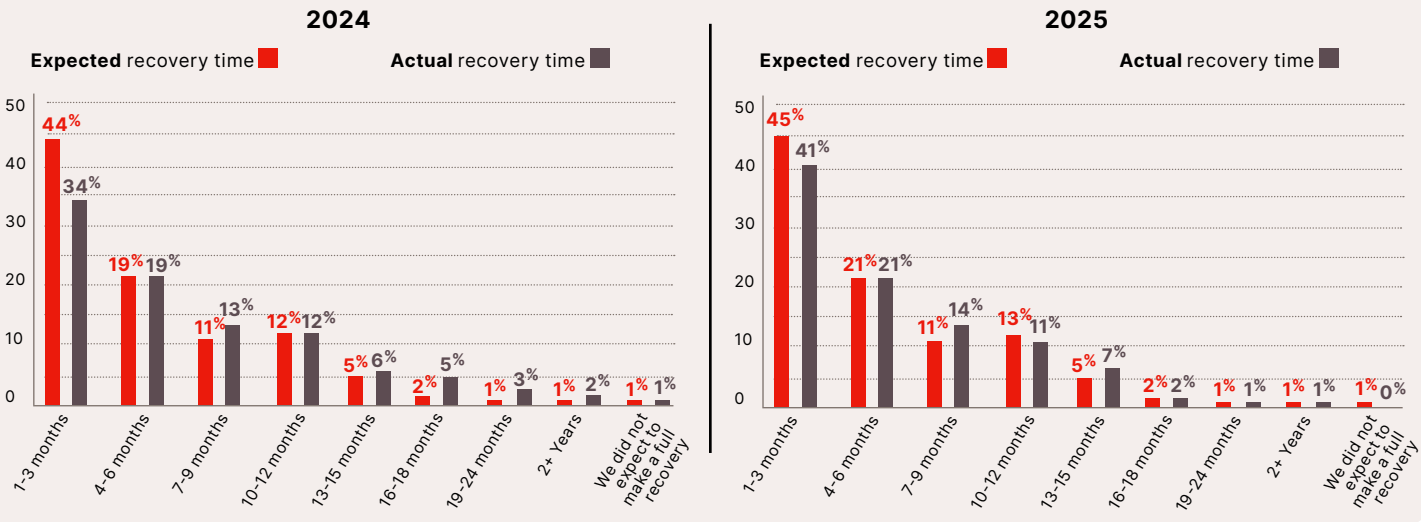

Investment priorities: where security dollars are really going

Ask security teams what keeps them up at night and data breaches naturally top the list at 45%. Social engineering comes in second at 40%. That makes sense, given that phishing emails remain the launch point for most successful attacks. Fancy technology still won't stop a harried employee from clicking a convincing link late on a Friday afternoon.

Ransomware (often launched using these social engineering attacks) ranks highly at 28%, after generative AI and a lack of technical skill. After years of headlines, that shouldn't surprise anyone, companies are still falling victim to it. Just because the playbook is well understood by now doesn't make it easier to stop.

These risks don't exist in isolation. Account takeover (19% cite this as a worry) also usually starts with phishing.

However, the third-party risks deserve particular attention. September's attack on the npm software repository showed how a single attack can infect hundreds of packages, trickling malware into countless user environments. When your security depends partly on someone else's, you're adding risk you can't directly control.

With mounting risks, it's no wonder that data protection and privacy lead investment priorities at 37%. The shift toward data protection reflects mounting regulatory pressure and the realization that data breaches carry consequences beyond immediate technical remediation. Compliance violations, customer lawsuits, and brand damage all stem from mishandled data.

Cyber insurance was the second most popular investment category at 34%, followed by API security at 33%. Cyber insurance's runner-up status is a sign of acceptance. Organizations realize that a perfect defense is impossible, so they're transferring risk. Can't prevent every breach? At least make sure you're not absorbing the entire financial hit when one lands.

With mounting risks, it's no wonder that data protection and privacy lead investment priorities at 37%.

The API security investments make sense when you consider that every new mobile app, third-party integration, and microservice creates another potential entry point. APIs used to be internal plumbing. Now they're exposed to the internet and attackers have noticed.

The real investment shift shows up in AI. The survey found that organizations are preparing for agentic infrastructure threats. Among those investing in this area, agentic discoverability tops the list at 56%, followed by API security at 55% and WAFs at 54%. These non-traditional security categories are responses to architectural patterns that barely existed two years ago.

Three quarters of companies are worried about DDoS attacks targeting AI agents, and over half admit their teams lack the AI specific expertise to handle these threats. Yet DDoS protection ranks just fifth as an investment priority at 30%. There's a disconnect between stated concerns and actual spending that suggests either misplaced priorities or constrained budgets forcing uncomfortable tradeoffs.

2 "Widespread Supply Chain Compromise Impacting Npm Ecosystem | CISA." Cybersecurity and Infrastructure Security Agency CISA, 23 Sept. 2025, www.cisa.gov/news-events/alerts/2025/09/23/widespread-supply-chain-compromise-impacting-npm-ecosystem

CISOs have more responsibility but hollow support

The CISO's chair is getting hotter. Nearly three quarters (73%) of organizations now say the CISO is ultimately responsible when breaches occur. Regulatory pressure is increasingly showing that security leaders can face personal liability for failures. For example, The EU's 2022 NIS2 Directive allows temporary suspensions of executives deemed incapable of fulfilling their cybersecurity responsibilities, along with damages to be levied against general managers and CEOs.

CISO involvement in incident response has also jumped, with 82% reporting active participation and 74% seeing increased CISO engagement over the past year. On paper, this looks like security getting the executive attention it deserves. In practice, the response has been less impressive.

When it comes to cybersecurity, everyone has a part to play

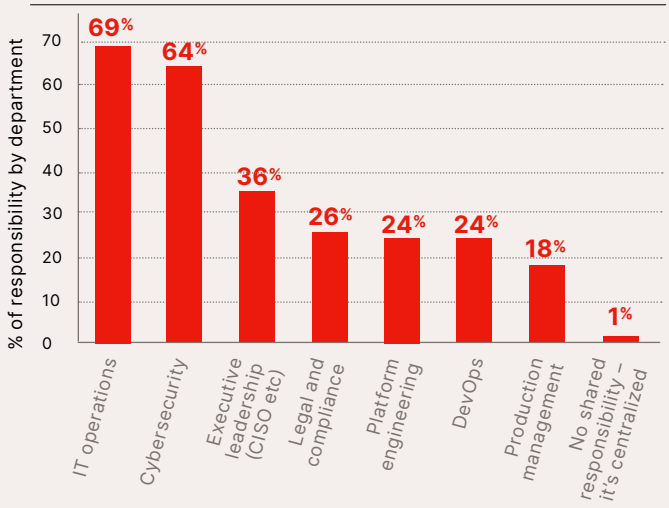

Policy changes miss the point

Most organizations (94%) made policy changes in response to growing CISO accountability. But dig into what those changes actually entail and the picture gets murky. Many measures are defensive, with 42% promising increased scrutiny of security disclosure documentation (in other words, reading the rules properly). Of respondents surveyed, 45% are now offering more legal support for cybersecurity staff.

These are colloquially known as 'CYA (cover your ass) policies'. "These measures are nice, but little more than self-preservation," says Erwin. "Those aren't actually improving your security posture."

One of the most common measures, cited by 44%, is finally giving the CISO a seat at the table for strategic decisions (that ties almost joint first with additional legal support and resources for cybersecurity teams, which each come in at 45%). That's at least giving them some kind of voice, but that alone isn't enough to stop the rot. Security leaders need the authority and resources to implement necessary security measures, not just the responsibility when things go wrong.

Nobody knows who's in charge

The confusion extends beyond technical vulnerabilities into organizational structure. Over half of AI-first businesses report a lack of clarity over who is responsible for incident response, compared to just 23% of non-AI-first organizations. Yet when breaches occur, blame flows uphill, with 79% of AI-first businesses reporting the CISO is ultimately held responsible versus 57% among traditional organizations.

As AI-first organizations create new risks faster than they can manage them, this issue will take center stage for security leaders. They will shoulder the blame as the inevitable breaches follow those risks.

Confusion doesn't just cause problems at the top. Our research shows that 43% of people within enterprises feel there's clarity around who's responsible for incidents – the same percentage feel there isn't (with the remainder on the fence).

3 PricewaterhouseCoopers. "What You Need to Know about NIS2 - PwC." PwC, 2023, www.pwc.de/en/cyber-security/european-nis2-directive-implications-for-businesses-and-institutions.html.

Solving the skills strain

The talent problem is getting worse, not better. More than half (53%) of security teams lack the AI specific expertise needed to respond to emerging threats. This isn't just about general cybersecurity skills anymore. The rapid adoption of AI infrastructure has created demand for specialized knowledge that the market can't supply.

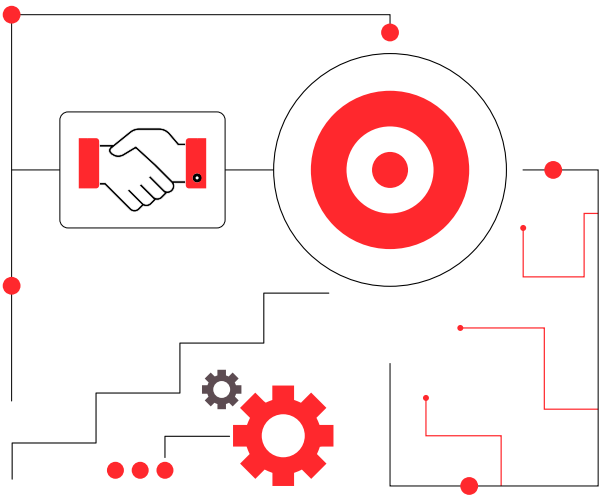

AI first organizations are especially hungry for talent. They're moving fast and building new systems. They're also discovering that traditional security expertise doesn't translate cleanly to protecting AI agents, managing agentic infrastructure, or defending against AI powered attacks. The skills shortage has become a bottleneck limiting how quickly they can secure their expanding technology footprints.

Fresh cybersecurity graduates face a steep learning curve. They must learn technical skills specific to a company's toolset and workflow, plus organizational cultural nuances. It takes substantial time and effort before a raw recruit becomes productive. As companies scale, this challenge intensifies, particularly when working in larger, constantly evolving environments.

Alternatives to external recruitment

Perhaps companies should look inwards instead. Several options deserve consideration. Upskilling existing staff for new responsibilities means they're already aligned with your culture and at least partly fluent in your specific systems and processes. These people understand the business context, which matters as much as technical capability when making security decisions under pressure.

Mentoring provides on-the-job training from experienced staff, cementing junior employees' skills and shaping them for success. It's slower than hiring someone with 10 years of experience, but those experienced candidates are increasingly hard to find and expensive to recruit.

Cross-functional collaboration between security and other teams like IT, compliance, support, and product development can create well-rounded employees with a strong sense of how security fits into other functions. There are opportunities for secondments here.

Sourcing talent from within, especially across different functions, carries several advantages. It promotes the idea that everyone is responsible for security. It also supports digital transformation efforts by embedding security expertise throughout the organization rather than concentrating it in a single team that becomes a bottleneck.

How threats vary by sector

The aggregate data tells one story, but drilling down into individual sectors reveals how unevenly the cybersecurity burden falls across different industries. Some face existential threats to their core business models, while others grapple with geopolitical risks or the operational chaos of securing sprawling digital estates. Here's how the threat landscape breaks down sector by sector.

Finance

Finance averages 54 breaches a year. Phishing still worries **39%** of organizations, and data breaches concern **42%**, but the killer statistic is **$442,232**. That's the average increased annual infrastructure cost from AI scraping alone, and it's the highest across all sectors. Perhaps that's why generative AI worries **41%** of financial institutions. Only tech companies are more anxious about it.

Media and entertainment

This sector breaks the mold entirely. It shoulders the lowest number of breaches across all sectors, at 24. However, content scraping is a business model threat, with one in five of these companies identifying it as a major concern, versus **5-16%** elsewhere. That shows up in the **51%** of media companies suffering elevated infrastructure costs from AI scraping, the **47%** encountering operational disruption, and the **39%** reporting customer experience issues – all of which are cross-sector highs. Just **11%** report no impact.

Government

Nation-state attacks are understandably the big issue for government respondents, worrying **21%** of security teams (markedly more than the percentage in commercial sectors). The percentage fretting about data breaches tops out at **52%**, the highest anywhere.

Retail

Retailers seem to get hammered from every direction. DDoS attacks are a concern for **25%** of them, (the highest rate across all sectors), while ransomware worries among retailers match those in the finance community at **32%**. And a sector-high **32%** skills gap means payment systems and customer data aren't properly defended.

Sectors may have differences that cause them to score differently in various aspects of cybersecurity, but one thing is universal: they all pay an AI tax now. Scraping costs run from nearly **$300k-$450k** annually across named sectors. The traditional threats such as identity attacks and data breaches haven't gone away, but AI has added mandatory new costs to running digital operations.

Financial organizations lost the most on average from their single biggest security incident in 2025

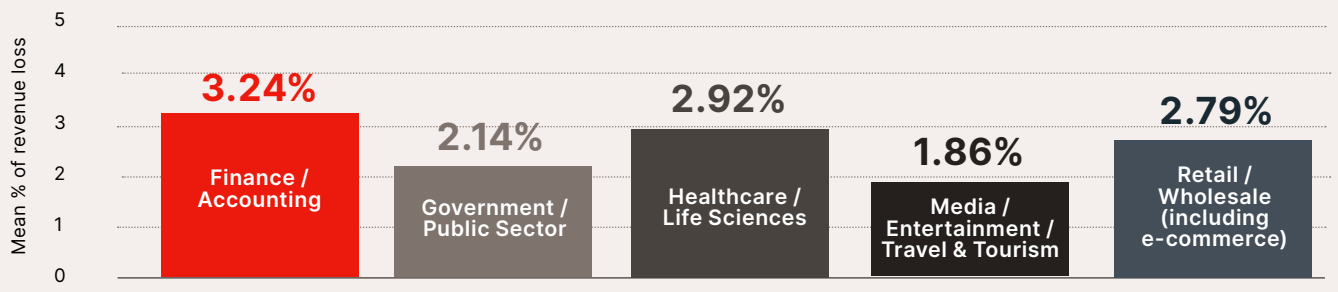

The path forward: security by design in an AI-accelerated world

With internal mistakes now causing as many breaches as external attacks and software bugs triggering 40% of incidents, you'd think this would spark a fundamental rethink of how organizations build software. But it hasn't.

> "I can't wait for someone to come to me for approval, because if that's happening, then I probably already failed."
>
> — Marshall Irwin, CISO at Fastly

Only 37% have shifted security responsibility to platform engineering or DevOps teams.

Participating in conversations before architecture decisions are made is crucial. "I can't wait for someone to come to me for approval, because if that's happening, then I probably already failed," Erwin says.

DevOps automation helps, but only with tactical problems. Your continuous integration/ continuous delivery (CI/CD) pipeline won't catch architectural mistakes like giving an AI agent excessive privileges that open your infrastructure to attack. Those require human judgment during design.

The automation ceiling

Automation is an important part of the security equation, because it helps to bridge the talent gap. However, even companies committed to automation hit limits. "The volume of potential vulnerabilities that you need to look at and get ahead of is high for just about any mature company in the tech space today," notes Erwin. Combine high volume with high false positive rates and automation only gets you so far.

Erwin advises automation for managing less serious issues and keeping expert humans for serious incidents. Let the machines handle routine vulnerabilities. Save human expertise for the tricky stuff that requires context and judgment.

Two problems, one solution

AI creates both external and internal security challenges. On one side sits the operational cost problem from external AI threats. AI scraping bleeds infrastructure budgets so badly that 64% of organizations now consider it a material cost center. That's because bots are crawling their sites, consuming bandwidth, degrading performance, and increasing their cloud bills (not to mention misappropriating valuable intellectual property).

On the other side sits the attack surface problem facing companies' own AI infrastructure. Agentic infrastructure and privileged AI tools create new vectors for attackers. These highly privileged tools can give attackers deep access to infrastructure if exploited.

Both problems converge at the same point: web applications and APIs. WAAP solutions defend both fronts. The same layer that throttles scrapers burning through your infrastructure budget also protects the APIs underpinning agentic systems. Web application firewalls that block layer-seven attacks work whether the target is traditional infrastructure or AI agents. Organizations investing in agentic security get this.

> "The volume of potential vulnerabilities that you need to look at and get ahead of is high for just about any mature company in the tech space today."
>
> — Marshall Irwin, CISO at Fastly

How to fix security before you start

AI-first businesses must learn the value of measured movement. So if they move fast and break things, it takes them longer to pick up the pieces. That's why AI organizations are increasing security spending yet feel more vulnerable.

It's also why security by design has moved from an aspirational goal to a survival requirement. The 81% who say resilience investments safely accelerated their innovation have already figured this out.

Security architecture built into systems from the beginning remove uncertainty, enabling teams to move faster with confidence. The alternative is what we're seeing now: organizations spending more, recovering slower, and wondering why buying more tools didn't fix anything. It's no coincidence that 72% of organizations prioritize speed-to-market over building resilience into systems.

AI is rewriting business operations at a pace that makes the economy's decade-long migration to the cloud look leisurely. Companies building that new infrastructure without security architects in the room from the first conversation are going to join the large community of AI-first organizations that saw AI directly exploited in their most recent breach.

There will be a demarcation line between those who embedded security into their AI strategy and those who bolted it on afterward. The good news is that you control which side you stand on.

Security benefits of implementing AI fro the start

Early implementation

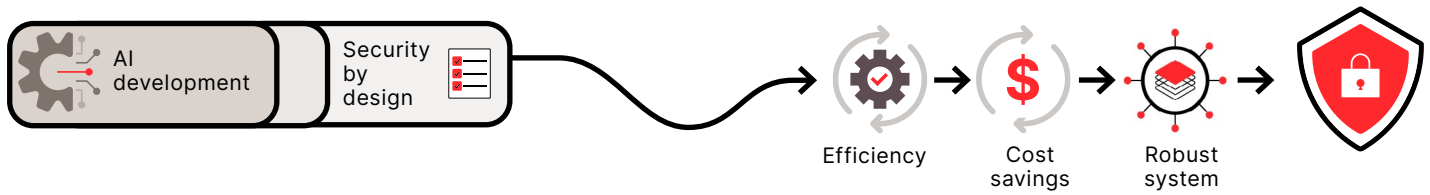

Mid-stream implementation

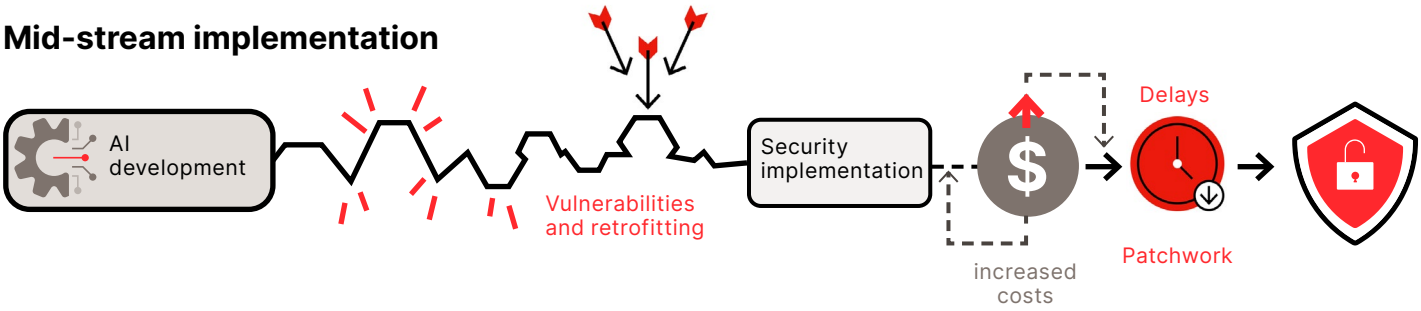

About the research

This research surveyed 2,000 key IT decision makers with an influence in cybersecurity, in large organizations spanning multiple industries across North, Central and South America, Europe, Asia-Pacific, and Japan. The interviews were conducted online by Sapio Research in Q4 2025 using an email invitation and an online survey. Results of any sample are subject to sampling variation.

The magnitude of the variation is measurable and is affected by the number of interviews and the level of the percentages expressing the results. In this particular study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 2.6 percentage points from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample.

About Sapio

Best new agency finalist, Sapio is adept at opinion polling (we have access to 80 million people internationally), focus groups, face-to-face interviews, telephone interviews, online research, desk research and statistical modelling to mention just a few techniques. We love B2B research and consultancy. Our business is based on partnership principles inspired by social enterprise.

About Fastly, Inc.

Fastly's powerful and programmable edge cloud platform helps the world's top brands deliver online experiences that are fast, safe, and engaging through edge compute, delivery, security, and observability offerings that improve site performance, enhance security, and empower innovation at global scale. Compared to other providers, Fastly's powerful, high-performance, and modern platform architecture empowers developers to deliver secure websites and apps with rapid time-to-market and demonstrated, industry-leading cost savings. Organizations around the world trust Fastly to help them upgrade the internet experience, including Reddit, Neiman Marcus, Universal Music Group, and SeatGeek. Learn more about Fastly at https://www.fastly.com, and follow us @fastly.