think report

Cost of a Data Breach Report 2026

The AI tipping point

IBM

01

04 Executive summary
05 What's new in the 2026 report
06 Key findings

02

08 Complete findings
10 Global highlights
18 Data security
22 Initial attack vectors and root causes
26 Identifying the breach
30 Data breach lifecycle
32 Breach recovery time
34 AI-driven attacks
38 Ransomware attacks
40 Breaches involving AI models and applications
48 AI oversight
50 Factors that increase or decrease breach costs
52 Security AI and automation
55 Agentic AI in security
56 Security investments
62 Securing non-human identities
63 Post-quantum cryptography

03

64 Recommendations to help reduce the cost of a data breach

04

66 Organization demographics
66 Geographic demographics
67 Industry demographics
67 Industry definitions

05

68 Research methodology
68 How we calculate the cost of a data breach
69 Data breach FAQs

06

71 About IBM and Ponemon Institute

Executive summary

Welcome to the 21st annual Cost of a Data Breach Report. Frontier AI models have radically shifted the cybersecurity threat landscape.

The announcement in April 2026 of a frontier model that managed to find thousands of high-severity vulnerabilities—including some in every major operating system and web browser—is a signal warning to security teams.

In the hands of attackers, these tools will collapse the time between vulnerability discovery and exploitation. Attackers are abandoning human speed for machine speed. They're already doing it with generative AI, which has lowered the time, cost and expertise needed to launch attacks, pushing organizations toward continuous business disruption.

This year's Cost of a Data Breach Report shows the cracks in the wall spreading. AI-driven attacks increased 56% over last year's study. The financial consequences were not minor. AI-driven attacks added an average of USD 1 million per breach as AI tools allow attackers to increase their velocity and scale. That speed is reshaping breach economics—and not in a good way.

The global average cost of a data breach climbed 12%, reaching nearly USD 5 million. That increase was largely driven by detection, escalation and lost business costs. AI and automation remain among the most effective ways to reduce breach impact. They help security teams move at machine speed and save an average of USD 1.93 million in costs. However, the adoption of these tools across the security lifecycle remains uneven.

We first noted this escalating AI arms race between attackers and defenders last year, when the global average breach cost dipped for the first time since the pandemic. Security teams seemed to be turning the tide. But this year's report shows them losing ground in some key metrics.

While 50% of breached organizations told us they've deployed agents in threat hunting, response and containment, only 18% applied them to vulnerability scanning and management—precisely where frontier capabilities set their sights. In response to new threats from frontier AI models, 85% of breached organizations said they plan to increase spending on security tools and governance across an array of solutions and services. That increase will be important as AI agents proliferate across organizations, requiring a focus on securing non-human identities (NHIs) in AI workflows. This year's research shows fewer than half of organizations securing those identities.

Meanwhile, attackers are not only using AI, they are targeting AI. Among the roughly one in five organizations that reported an AI-related breach, we found security incidents were less about model selection than model and environmental security. Similar security incident rates occurred across open-source and third-party vendor deployments. The root causes of these incidents were often structural: compromise of connected APIs, applications and cloud misconfigurations, indicating governance failures—not model risk.

On the plus side, organizations are recognizing the need to strengthen security for AI. More than half say they plan to invest in AI security and governance tools post-breach, an 88% increase from last year and a reaction to concerns over frontier AI model threats. That increased focus on security is good news. AI sovereignty—maintaining control over where AI systems run, how data is processed and who has access—has become a critical security requirement. Without it, organizations expand their attack surface and create dependencies that can delay detection and containment.

This year's research—conducted by Ponemon Institute and sponsored, analyzed and published by IBM—studied 602 organizations impacted by data breaches from March 2025 through February 2026 (the years mentioned in this report refer to the year the report was published, not necessarily the year of the breach). Together, we looked at organizations across 17 industries, in 16 countries and regions, and breaches that ranged from 2,590 to 115,380 compromised records.

Ponemon interviewed 3,558 security and C-suite business leaders with firsthand knowledge of the breach incidents at their organizations. These leaders included CEOs, CISOs, heads of operations, controllers or heads of finance, IT practitioners, business unit leaders and general managers, and risk management and cybersecurity practitioners.

The findings and analysis in this year's report can help business, technology, and security and risk leaders strengthen their defenses, inform resource allocation and drive innovation. To make progress against new frontier AI-enabled threats, it's helpful for leaders to understand the starting point of security teams. What are security operations center (SOC) teams doing right, and what areas need to be improved? This year's Cost of a Data Breach Report paints that picture in vivid detail.

What's new in the 2026 report

As always, the Cost of a Data Breach Report reflects new technologies, emerging tactics and recent events. For the first time, this year's research explored the following:

- Areas of agentic AI use in the SOC
- Types of controls used to secure NHIs
- Prevalence of encryption technology for securing data
- Prevalence of post-quantum cryptography projects
- Types of controls to secure cryptography
- Breach cost reduction from using secret lifecycle management tools

Key findings

USD 4.99M

The global average cost of a data breach

The global average breach cost climbed 12% over last year, reaching a record USD 4.99 million. Put another way, that's an USD 1,100 per hour cost. Two cost categories—detection and escalation and lost business—made up the majority (63%) of costs in this year's report. Those costs include everything from crisis management to disrupted operations and customer churn. Average breach costs in the US—where regulatory fines and other business costs are higher—exceeded all other countries and regions in the study with a record USD 11.5 million. That's more than twice the global average and up 13% over last year.

USD 6M

The average cost for AI model inversion attacks

Attackers aren't only using AI as a tool, they're targeting the AI models and applications organizations increasingly rely on. The most expensive types of security incidents involving an organization's AI model or applications were inversion and prompt injection attacks, which led to average losses of USD 6.07 million and USD 5.89 million respectively. Those types of attacks reflect the growing difficulty of protecting training data, controlling model behavior and safeguarding sensitive outputs. Unlike traditional exploits that compromise systems, these behavioral attacks undermine how AI models reason and respond, expanding remediation beyond technical recovery into trust and governance.

56%

Increase in AI-generated attacks

Attackers are increasingly weaponizing AI, using it to automate the attack lifecycle, scale operations and refine the sophistication of their attacks. This year's research shows the impact: more than one in four organizations that experienced a malicious attack reported it was AI-driven, a 56% increase over last year. AI deepfake impersonation and AI-enabled malware drove the highest volume of those incidents. The financial consequences: AI-driven attacks added an average of USD 1 million per breach.

92%

Share of organizations that reported an AI-related breach and lacked proper AI access controls

Among organizations that experienced an AI-related breach, the vast majority (92%) lacked proper AI access controls. That's despite organizations reporting identity and access management was one of the most effective cost reducers across breaches overall. As AI systems scale and embed deeper into organizational workflows, identity and access controls have failed to keep pace. In this year's report, only 40% of organizations reported using access controls on AI models and data. The outcome is predictable: expanded attack paths, higher financial impact and incidents driven by basic enforcement gaps that don't even require attacker sophistication.

41%

Share of ransomware attacks that included threats to brand reputation

Ransomware incidents among breached organizations rose slightly this year over last year (39% versus 34%), as attackers used AI to generate malware and expanded their methods of extortion. While disrupting operations through encrypting remains a key tactic (23%), attackers are shifting to higher-impact pressure methods. Threatening brand reputation is now the most common tactic, cited by 41% of the breached organizations that experienced a ransomware incident. This shift reflects a move away from purely technical disruption toward multilayered extortion strategies that target trust, public perception and long-term business impact.

USD 6.29M

Average cost of a breach in the financial services industry

Critical sectors face the brunt of AI-driven attacks, accounting for 62% of all AI-driven breaches studied. The two most targeted sectors were financial services and energy, where breach costs averaged USD 6.29 million and USD 5.2 million respectively. This targeting is a concerning trend. The concentration of AI-driven attacks across these two interdependent sectors creates a systemic risk. Disruptions in one organization can cascade into broader impacts across consumer finances, economic systems and power grids. On the plus side, both sectors are experienced at identifying and containing breaches, with approximately four weeks shorter average breach times than the global average.

85%

Share of organizations that plan to increase security spending in response to frontier AI model threats

Originally, 64% of organizations said they would increase security spending after experiencing a breach. Upon learning of new threats from frontier AI models, that figure increased dramatically to 85%. The new threats have also caused organizations to rethink the use of AI agents in security. Three quarters said they will deploy agents at higher rates in alert triage, vulnerability management, and scans and penetration testing.

USD 1.93M

Cost savings from extensive use of AI in security operations

Security teams extensively using AI and automation shortened their breach times by 65 days and lowered their average costs by USD 1.93 million compared to those that didn't use these solutions. These cost savings reinforce the idea that modern security operations must operate at machine speed. However, only about one-third (36%) of breached organizations said they extensively used these tools across the security lifecycle in prevention, detection, investigation and response. Their use in prevention lagged behind detection and investigation workflows.

50%

Share of organizations deploying AI agents in their SOC

While AI and automation remain among the most effective ways to reduce breach impacts, their adoption across the security lifecycle remains uneven. They're most commonly used for detection and investigation while use in prevention lags. Agentic AI adoption reflects the same uneven pattern. Among the half of breached organizations that said they deploy agents in security, more than 50% are focused on threat hunting, response and containment. Only 18% applied AI agents to vulnerability scanning and management, an application that could help reduce exposure in the first place and potentially blunt the impact of frontier AI threats.

Complete findings

The complete findings from this year's research address 16 themes, presented in the following order:

- Global highlights
- Data security
- Initial attack vectors and root causes
- Identifying the breach
- Data breach lifecycle
- Breach recovery time
- AI-driven attacks
- Ransomware attacks
- Breaches involving AI models and applications
- AI oversight
- Factors that increase or decrease breach costs
- Security AI and automation
- Agentic AI in security
- Security investments
- Securing non-human identities
- Post-quantum cryptography

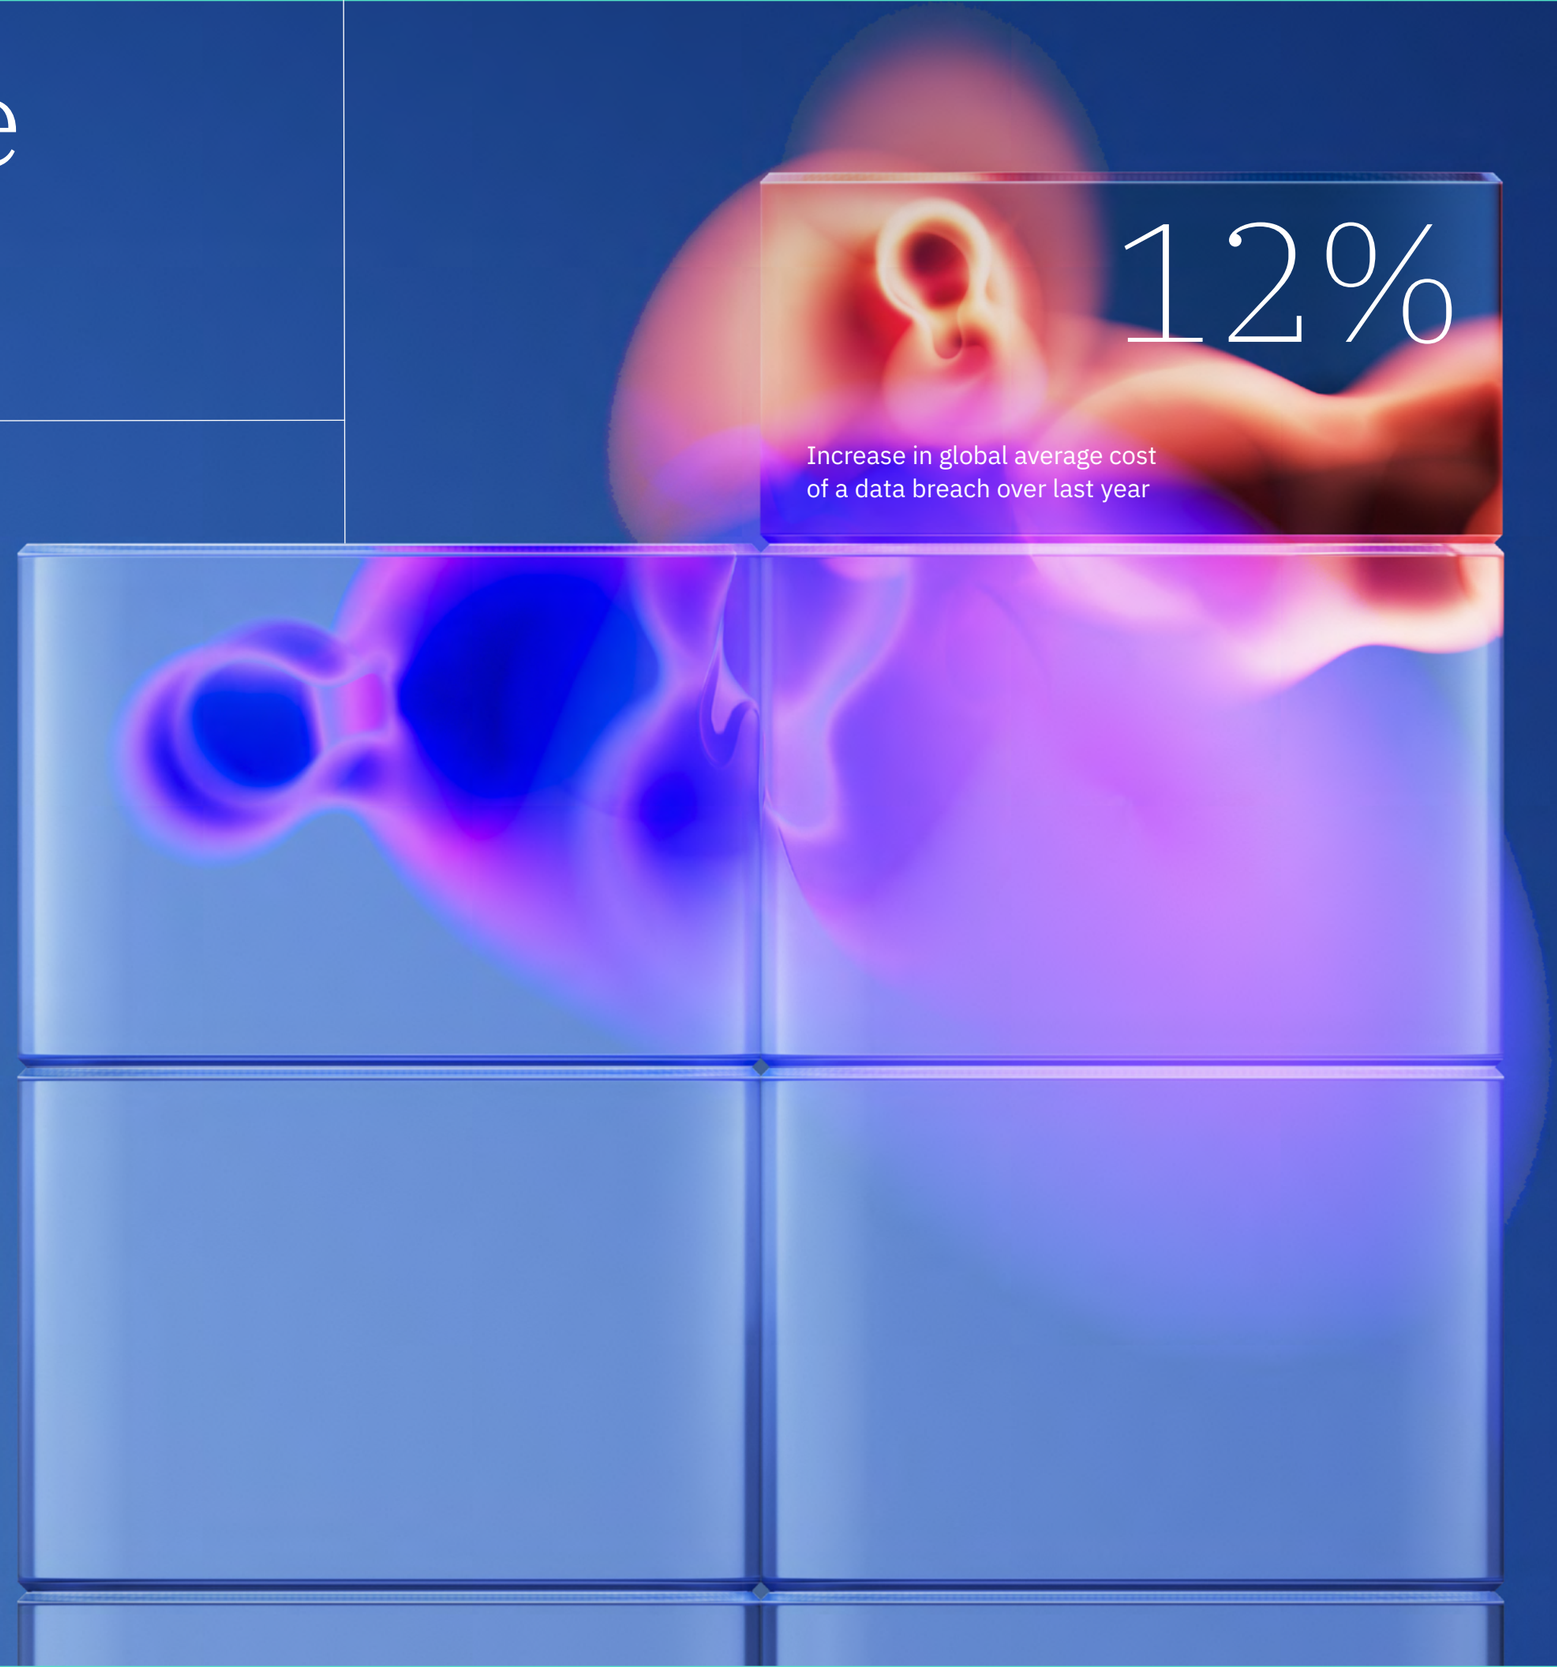

Global highlights

AI-driven attacks are re-shaping breach economics. They're getting faster and cheaper to launch, while breaches are getting more expensive to find and fix.

This year's average breach costs rose across all regions and industries, except healthcare. The United States again had the highest average costs among all regions. Healthcare continues to top the list of costliest industries for breaches. The following section provides a look at these and other issues.

Global costs hit a record high

The global average cost of a data breach rose to an all-time high this year, reaching USD 4.99 million. That figure represents a 12% spike over last year—when the average cost dipped 9%—and continues an upward trajectory since the pandemic. This increase was driven largely by detection, escalation and lost business costs, which includes costs factors related to disrupted operations and customer churn. See Figure 1.

Figure 1.
Measured in USD millions

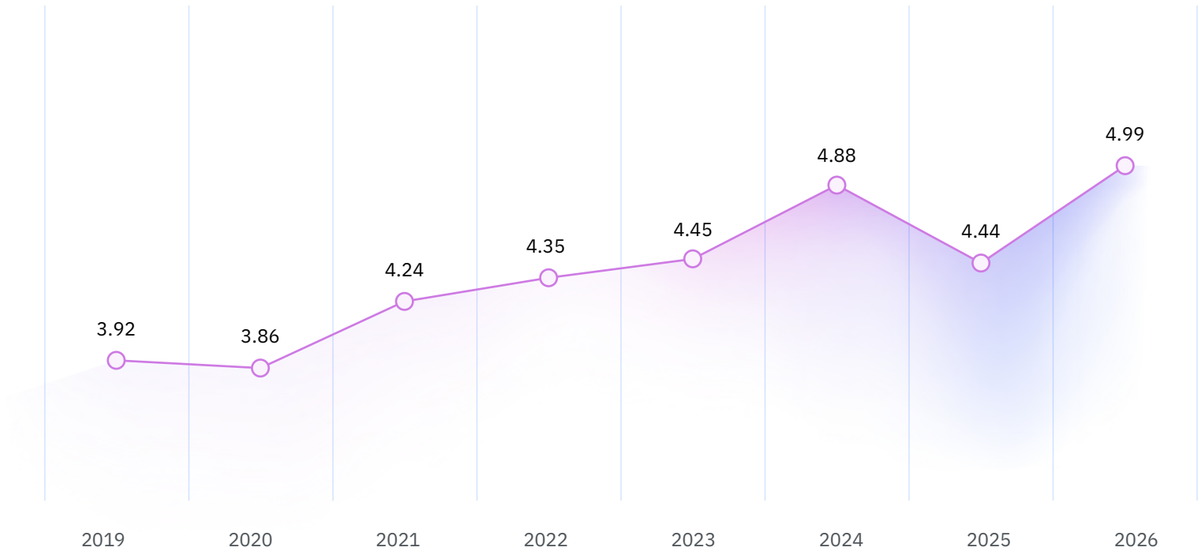

The United States breaks a breach cost record

Average breach costs in the US reached a record USD 11.5 million, an 11% increase over last year and nearly double the global average. Historically, higher business costs and higher regulatory fines in the US have contributed to this cost difference. Meanwhile, South Africa saw the largest percentage increase in average breach costs (22%) even though the actual cost (USD 3.04 million) was lower than in other regions. Average breach costs also rose in Germany (up 18% to USD 4.93 million) and Benelux (up 16% to USD 7.37 million). See Figure 2.

Figure 2.
Measured in USD millions

#	Country	2026	2025
1	United States	11.50	10.22
2	Middle East	8.00	7.29
3	Benelux	7.37	6.24
4	Canada	5.20	4.84
5	Germany	4.93	4.03
6	LATAM	4.65	3.81
7	United Kingdom	4.17	4.14
8	Italy	4.12	3.44

#	Country	2026	2025
9	ASEAN	4.12	3.67
10	France	4.05	3.73
11	Japan	4.01	3.65
12	South Africa	3.04	2.37
13	South Korea	3.03	2.84
14	Australia	2.96	2.55
15	India	2.79	2.51
16	Brazil	1.41	1.22

Detection, escalation, and lost business costs increased
Cost increases in two categories—detection and escalation and lost business—drove a majority (63%) of this year's record average breach cost. They each rose 11.5% over last year and together accounted for USD 3.18 million of the USD 4.99 million global average. The remaining two cost categories also rose. Post-breach response, which includes regulatory fines, legal expenditures, and credit monitoring for customers affected by a breach, rose the most (15%). See Figure 3.

Figure 3.
Average cost of a data breach in four components; measured in USD millions

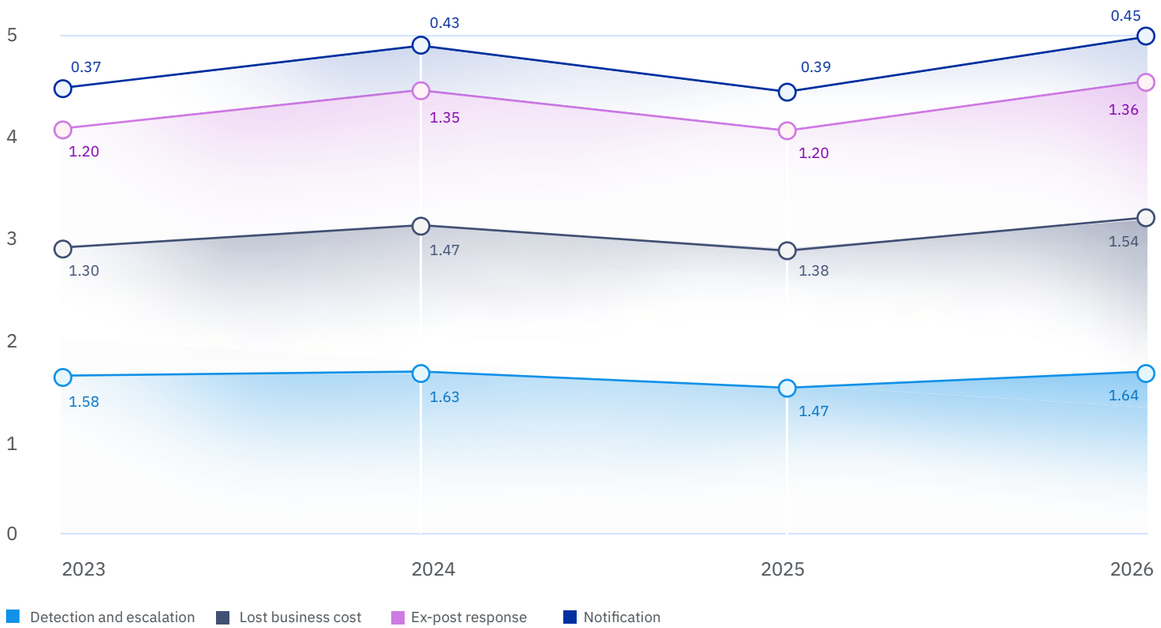

Figure 4.
Average total cost of a data breach by industry in USD millions

Healthcare remained the most expensive industry for breaches
For the 13th consecutive year, the healthcare industry recorded the highest average breach cost (USD 6.64 million) among all industries. Still, that figure is down 10.5% from USD 7.42 million last year. Attackers continue to value and target the industry's patient PII, which can be used for identity theft, insurance fraud and other financial crimes. Average breach costs also increased in all industries in this study. Media-related industries saw the biggest jumps, with average communications industry breach costs increasing 20% to USD 4.71 million, and entertainment costs rising 18% to USD 5.38 million. See Figure 4.

Breach response times rose
The mean time organizations took to identify and contain a breach rose to 247 days, a slight 2.5% uptick that reversed a five-year decline. In that time, security teams using AI and automation have steadily improved their mean time to identify (MTTI) and mean time to contain (MTTC) attacks. However, new threats from AI-driven attacks are challenging even the quickest response times and rewiring the security defense clock. See Figure 5.

Figure 5.
Measured in days

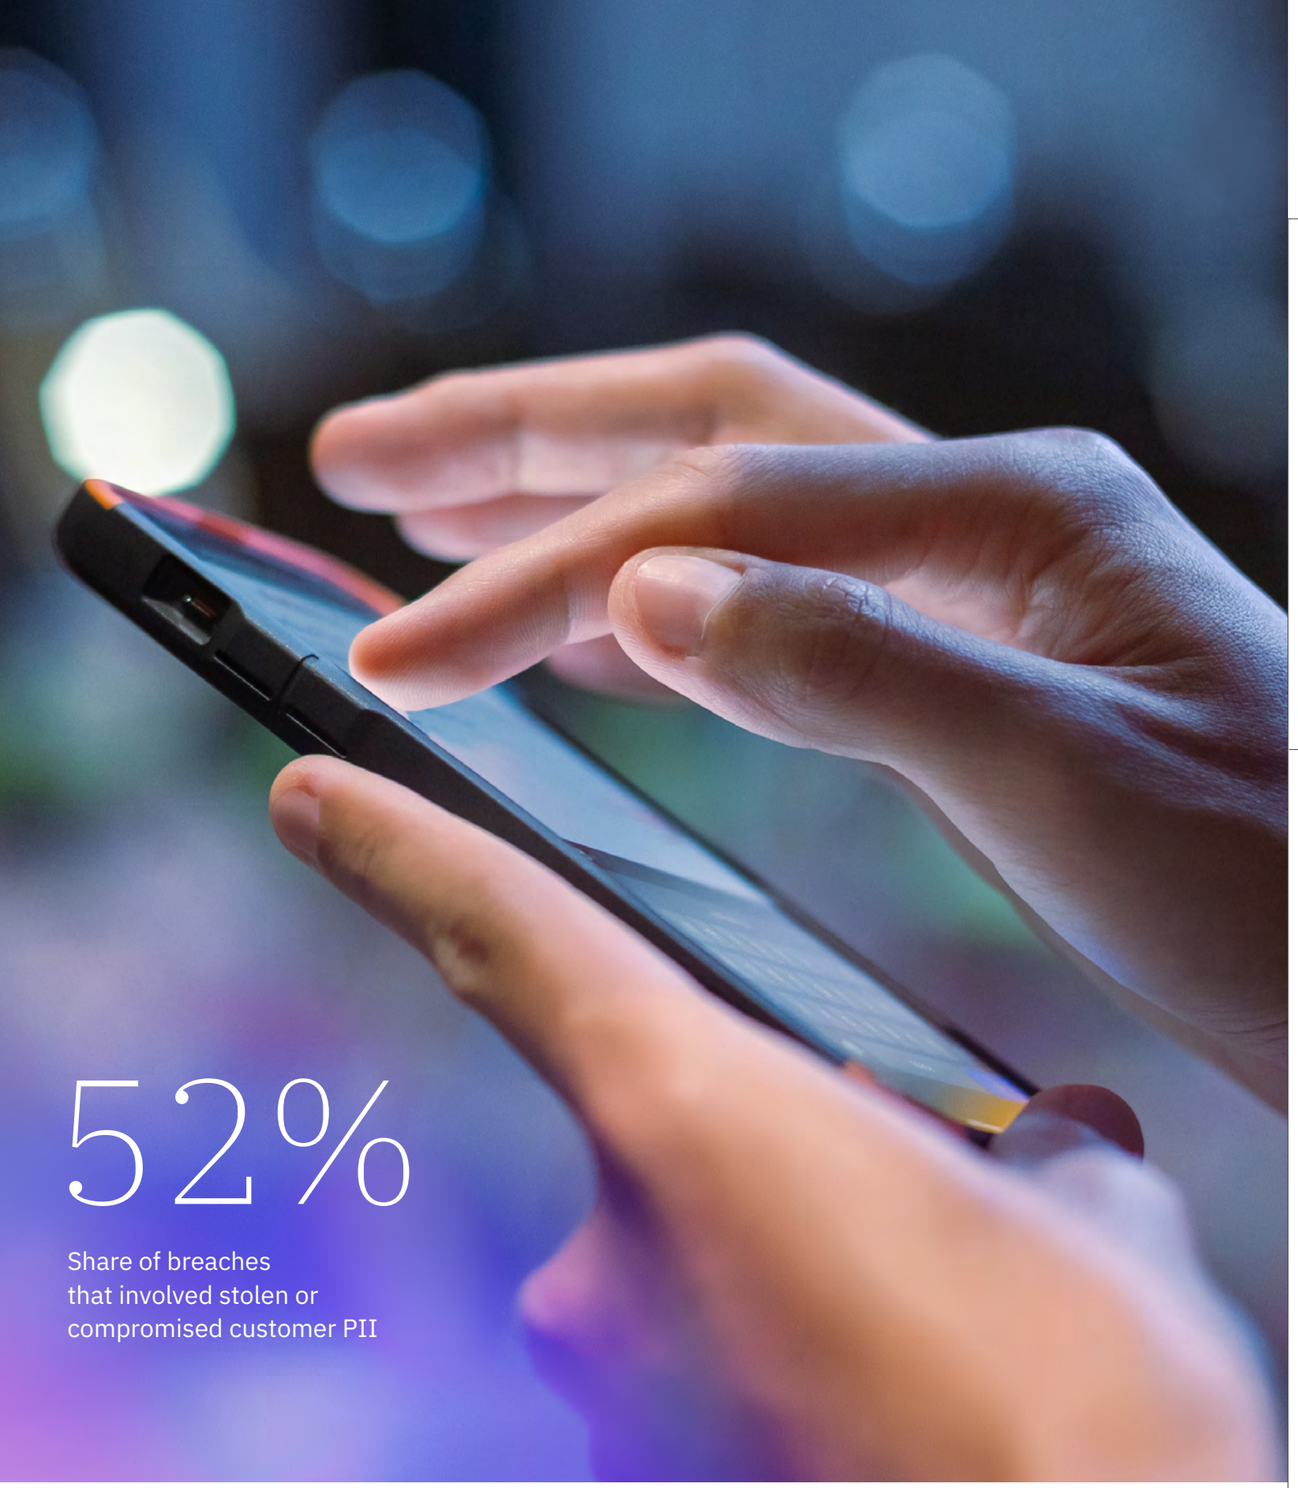

Most breaches targeted customer PII

Attackers targeted customer PII over other types of data by a wide margin. At 52%, it was the most stolen or compromised data type this year, costing USD 192 per record on average. Customer PII can include tax identity (ID) numbers, emails and home addresses, and can be used in identity theft and credit card fraud. Employee PII was targeted in 35% of attacks and cost as much as USD 188 per record on average. Intellectual property (IP) was stolen or compromised in nearly a third of data breaches (32%). It was the costliest type of breached data at an average of USD 196 per record. See Figure 6.

Figure 6.
Types and share of data stolen or compromised among breached organizations; more than one response permitted; measured in USD

Data security

The storage of data and its handling have become critical security and compliance issues as organizations and regulators wrestle with growing concerns over data sovereignty.

Data sovereignty is the principle that nation-states, local governments and regional alliances, such as the EU, should have control over how and where their residents' data is stored and processed. Where and how data is stored can affect its resilience to data breaches. This year, we found a greater share of breaches involved data stored on premises than in any other location. We also found use of data encryption—critical for preventing theft and protecting privacy—was missing from most breached organizations.

Figure 7.
Types and share of data stolen or compromised by location; measured in USD millions

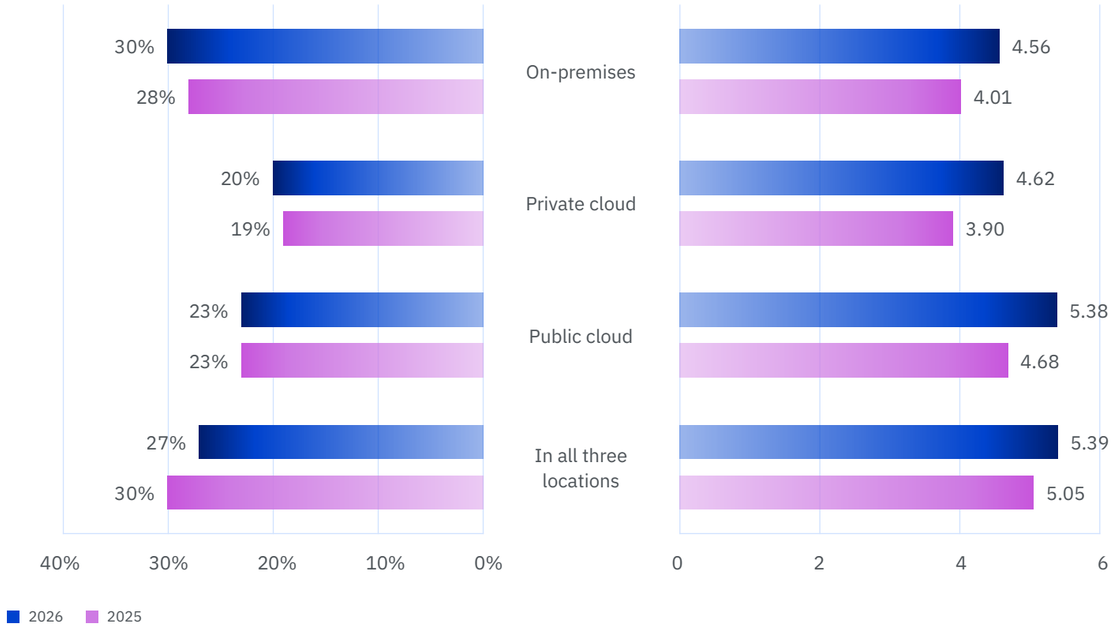

Breaches involving on premises data storage grew
In most breaches this year, the breached data was stored on premises (30%). That share has crept up from 28% last year and from 20% in 2024. Although many organizations believe on-prem storage is inherently safer, it remains a target because it places the sole responsibility for patching, physical security and access management on the organization's internal IT staff. It's where organizations tend to keep their most valuable assets and also where attackers can cause the greatest disruption. However, these breach locations also showed the lowest average breach costs (USD 4.56 million) among all locations, though that figure is 12% higher than last year. The most expensive storage locations were public cloud (USD 5.38 million) and data stored across all three locations (USD 5.39 million). See Figure 7.

Response and containment times varied across data storage locations
For the second straight year, breached data stored across multiple locations took the longest to identify and contain (256 days). That response time indicates the complexity and uncertainty of such breaches, but it's also an improvement over last year (277 days). Additionally, resolution times rose for some categories and fell for others. Private cloud breaches were the quickest to resolve (240 days), while on-premises breaches took longer (244 days) compared to last year's low (217 days). See Figure 8.

Figure 8.
Measured in days

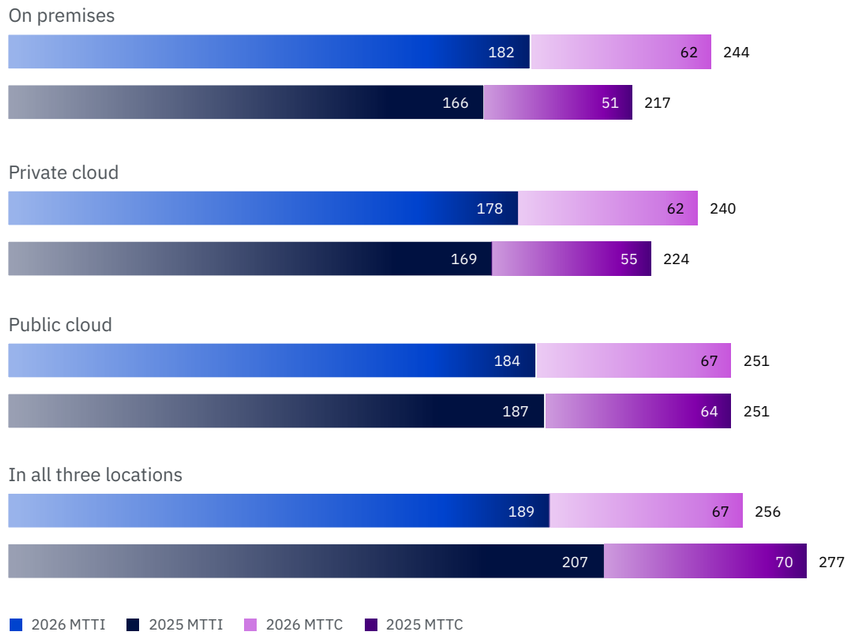

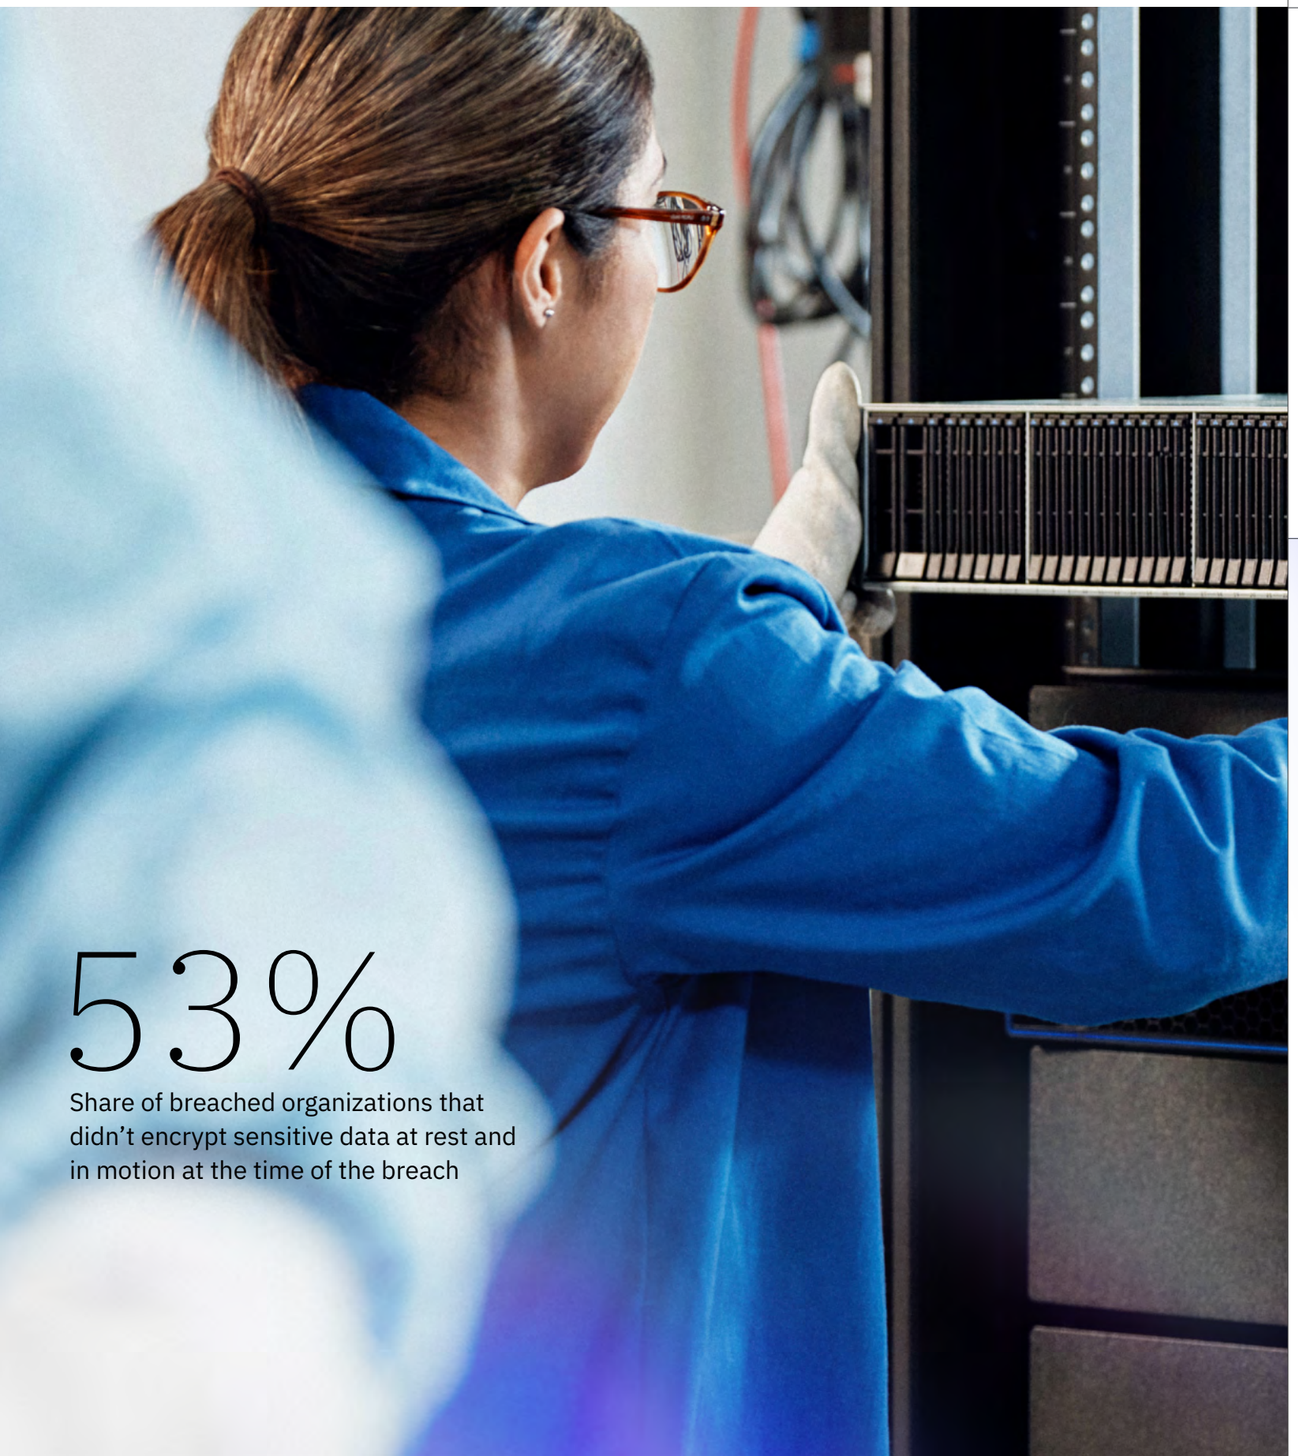

53%

Share of breached organizations that didn’t encrypt sensitive data at rest and in motion at the time of the breach

Most breached organizations left sensitive data unencrypted
Among breached organizations, 53% didn’t encrypt sensitive data at rest and in motion at the time of the breach. Another 10% of organizations said they weren’t sure if the data was encrypted. Unencrypted sensitive data enables attackers to gain full, immediate access to genetic, biometric, identity and health-related information. Under certain regulations, exposing this data is considered a severe risk to an individual’s and organization’s rights, and a violation of the foundational components of digital sovereignty. See Figure 9.

Figure 9.
Share of data encryption among breached organizations

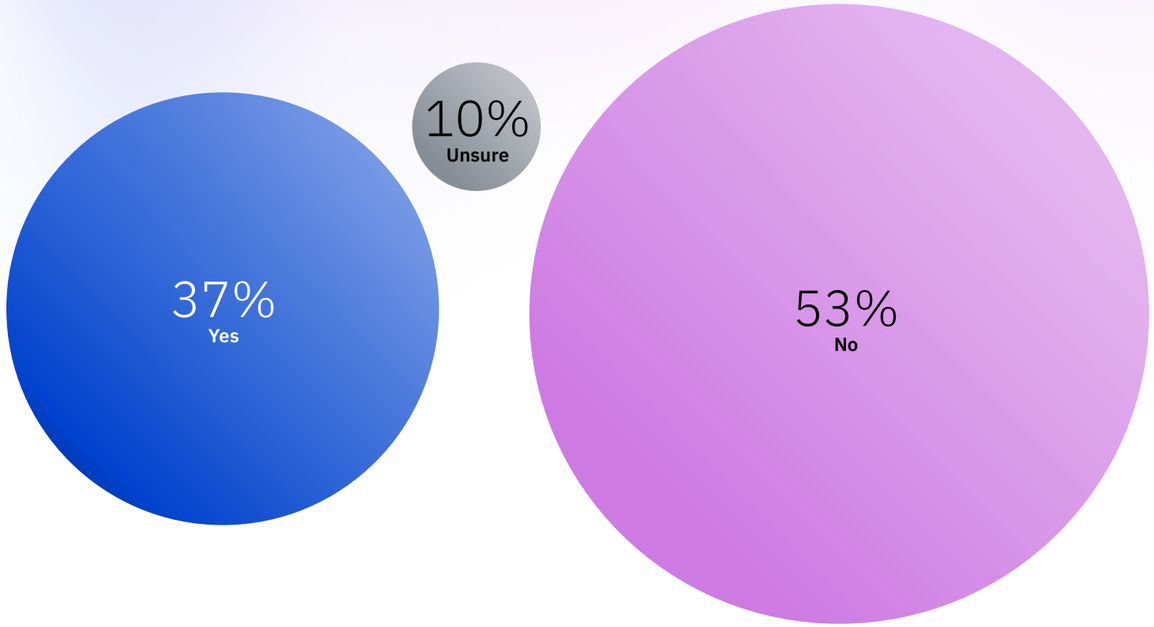

Initial attack vectors and root causes

For the fourth year in a row, phishing was the top attack vector into breached organizations. Supply chain compromise came second, followed equally by abusing valid accounts, drive-by compromise attacks and social engineering. Attacks that used data replication on removable media took the longest to resolve, closely followed by supply chain compromise.

Phishing topped initial attack vectors and led to the costliest breaches

Attackers used impersonation for conducting high-damage data breaches. Voice and SMS phishing, used in 17% of attacks, led to the highest average breach costs among attack vectors, USD 5.29 million. Social engineering, such as impersonating help desk staff, was used in 13% of attacks and led to average breach costs of USD 5.23 million. Valid account abuse was slightly less costly at USD 5.07 million, suggesting that phishing and social engineering attacks are targeting credentials with access to prized data. In nearly 10% of data breaches, attackers illicitly replicated data on removable media, showing that even in the age of cloud computing and AI, it's important not to overlook these decades-old threat vectors. Sometimes data still walks out the door on a thumb drive. See Figure 10.

Figure 10.
Measured in USD millions

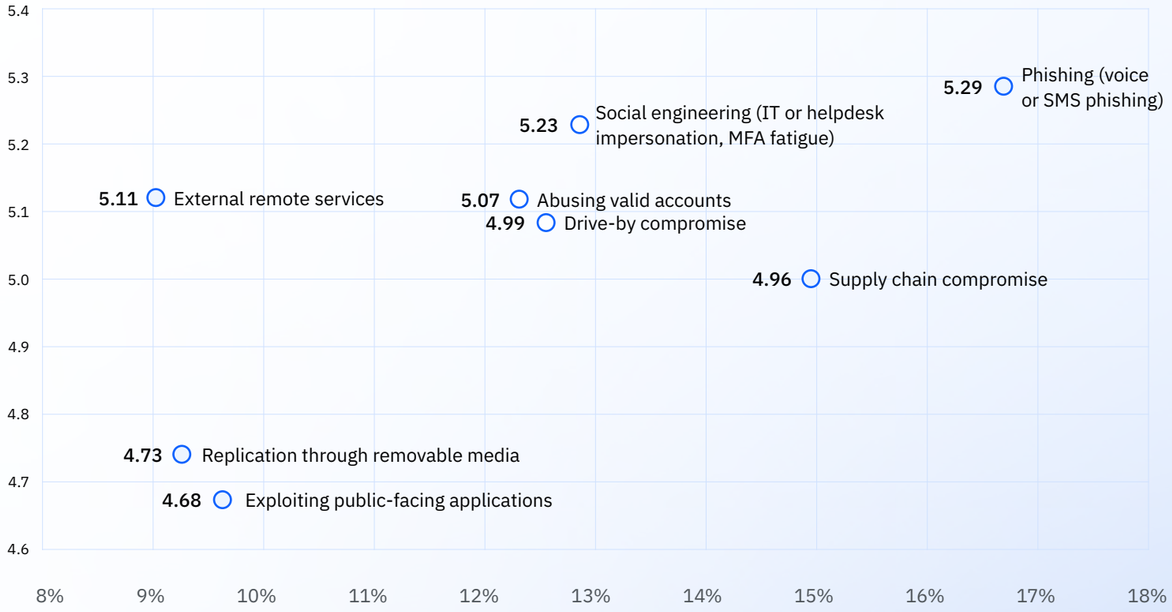

Figure 11.
Share of all breached organizations by industry sector

Malicious attacks remained the most common root cause of data breaches

Malicious or criminal attacks accounted for 55% of all data breaches, up almost 8% from last year. These attacks far more common than breaches caused by IT failures (22%) or human error (23%). Root causes, however, varied by industry. Malicious or criminal attacks were responsible for a little more than half (61%) of breaches in the retail and transportation industries and roughly half of breaches in the healthcare and finance sectors—industries known for storing valuable PII. However, in the entertainment industry, these attack types accounted for only 45% of breaches. See Figures 11 and 12.

Figure 12.
Share of all breached organizations

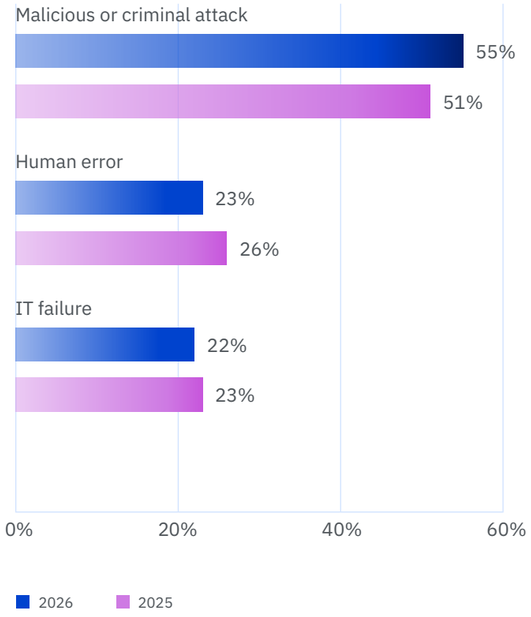

Removable media and supply chain compromises took longest to identify and contain

Breached organizations struggled with identifying and containing breaches that involved copying data onto removable media. These breaches don't show up in malware scans or inbound traffic. They might be noticed only when the stolen data appears on the dark web. These breaches lasted 258 days on average. Supply chain compromises also lasted 258 days. These attacks take advantage of trusted third-party software or other resources to give attackers access to an organization's data. Breaches involving drive-by compromise attacks were identified most quickly (234 days), possibly because scanning software was able to detect the malware that had been downloaded when an employee visited an infected website. See Figure 13.

Figure 13.
Measured in days

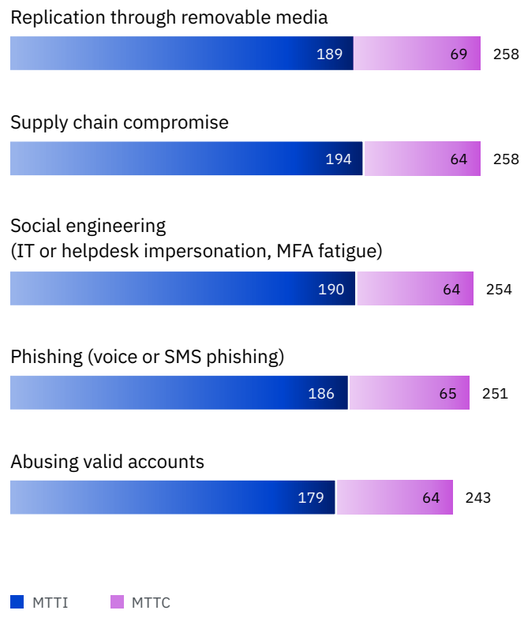

Identifying the breach

Breach costs rise or fall depending on who detects them and when. For the second year in a row, in-house security teams led the way in identifying data breaches. Researchers looked at the prevalence of breach disclosures by outsiders, such as benign third parties, security researchers, law enforcement and consultants, and by the attackers themselves. They then examined the costs associated with each type of breach identification.

Internal IT or IT security teams identified most breaches
Internal teams identified 38% of data breaches, many likely aided by AI and automation tools. Managed security service providers (MSSPs), who also have access to these tools, identified almost as many (31%). Third parties, such as partners, consultants or law enforcement, identified the fewest among this group (14%). See Figure 14.

Figure 14.
Share of breaches by who identified them

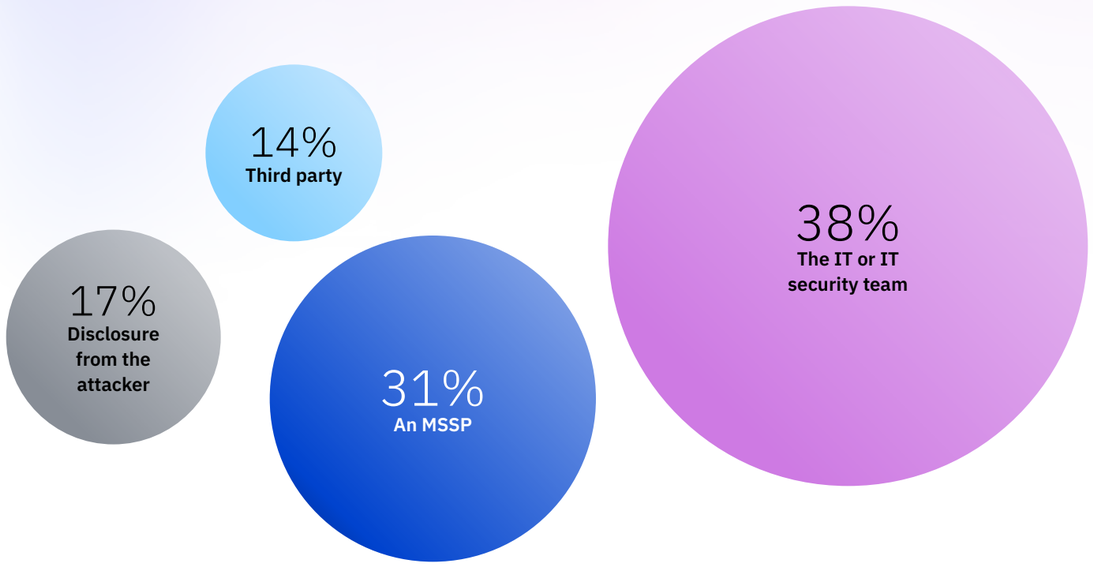

Attacker-disclosed breaches were the costliest
When breached organizations were unaware of a breach until attackers notified them, the average cost reached USD 5.12 million. By the time of disclosure, attackers can seize control of significant PII or gain leverage over operational systems they can disrupt. When internal teams identified the breach, the average cost dropped modestly to USD 5.01 million, despite faster response times. The US, which provides the largest sample size in the research and spends the most in IT security, drove up the average costs. Meanwhile, if a breach was identified by an MSSP—which can have a variety of advanced tools and security analysts—the average cost was even lower at USD 4.86 million, 2.6% below the global average. See Figure 15.

Figure 15.
Measured in USD millions

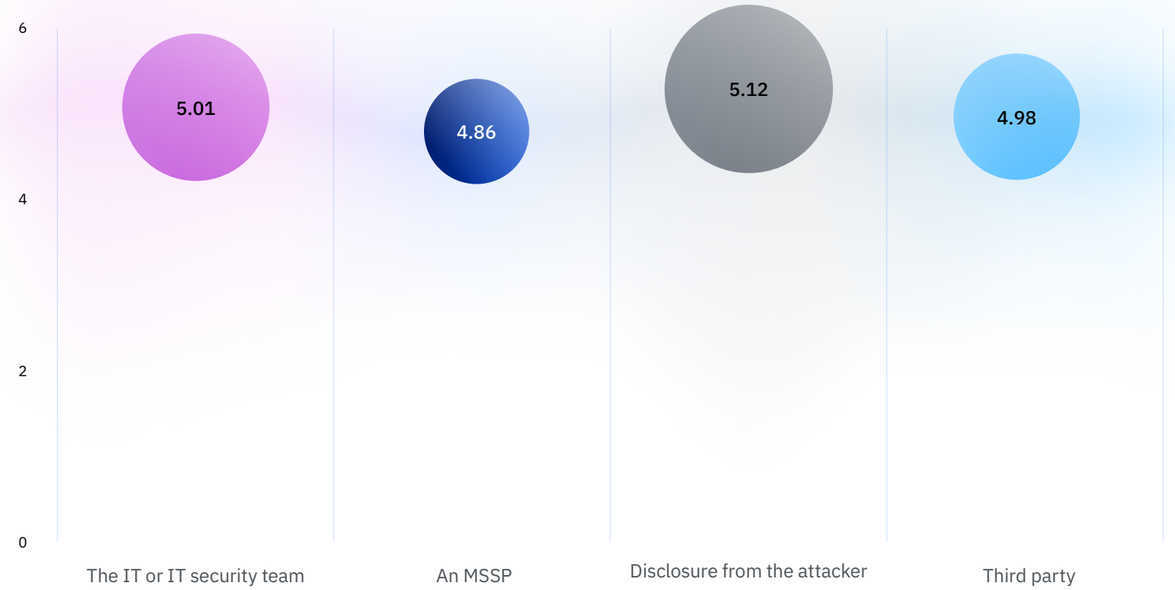

Internal teams identified breaches fastest
Internal IT or IT security teams identified and contained breaches in just 209 days on average—15% faster than the global average of 247 days. When MSSPs identified and contained breaches, the average timeline of 230 days was still faster than the global average. However, when third parties—whether business partner, consultant or law enforcement—identified and contained breaches, the average timeline stretched to 280 days. See Figure 16.

Figure 16.
Measured in days

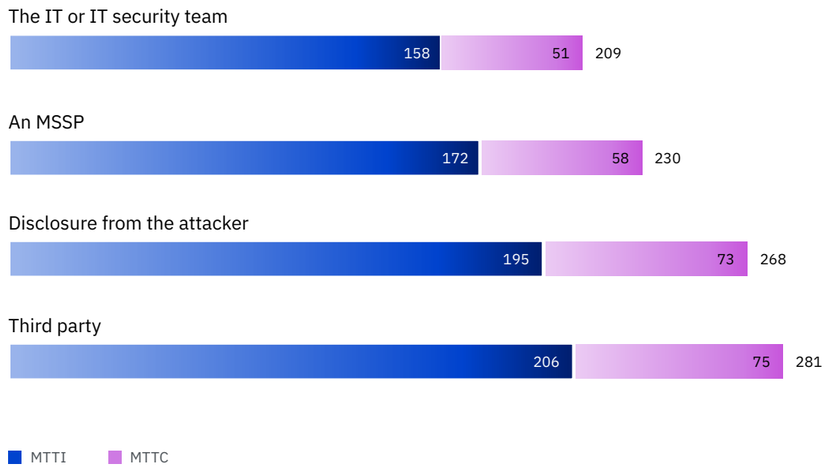

Data breach lifecycle

When an attacker breaches an organization, costs go up by the day. Each year, researchers analyze the average costs of the complete breach lifecycle—the total average number of days to identify and contain the breach—by breaking them into two categories: those that took less than 200 days and those that exceeded 200 days. The difference in cost between longer and shorter breaches this year was significant.

12%

Increase in average costs for data breach lifecycles longer than 200 days

Longer breach lifecycles led to higher costs
Data breach lifecycles longer than 200 days saw a 12% increase in average costs to USD 5.65 from USD 5.01 last year. Shorter-lived data breaches were far less costly, at USD 4.32 million, but marked an increase from last year's USD 3.87 million—an 11% jump. See Figure 17.

Figure 17.
Measured in USD millions

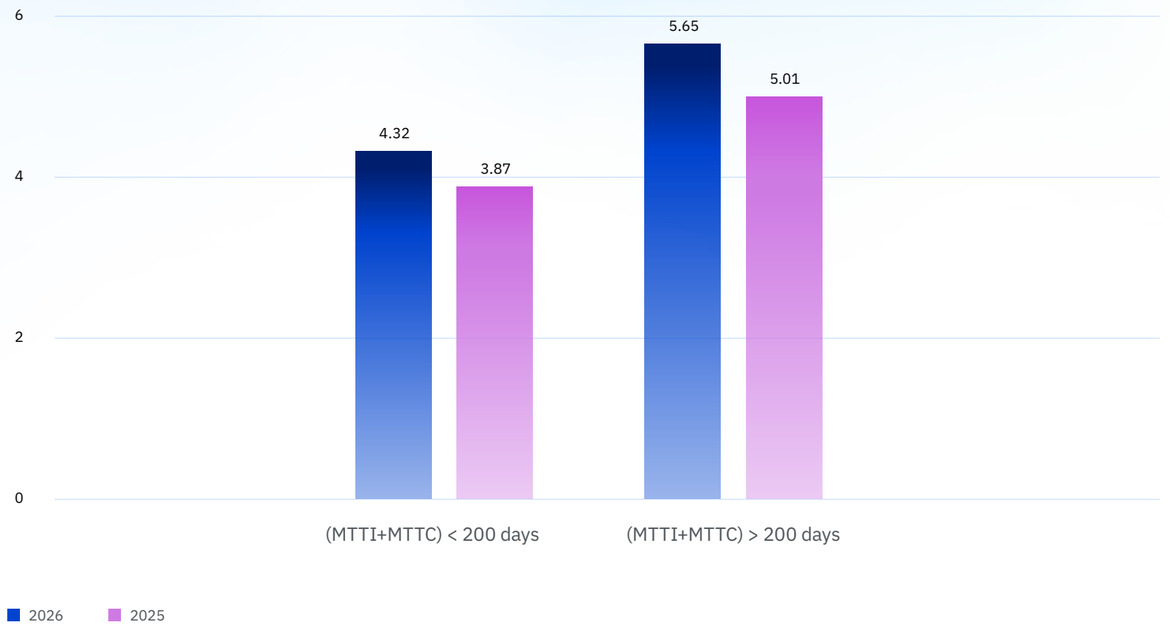

Breach recovery time

Recovery from a breach can continue after containment. In this study, recovery means:

- Business operations are back to normal in areas affected by the breach.
- Organizations have met compliance obligations, such as paying fines.
- Customer confidence and employee trust have been restored.
- Organizations have put controls, technologies and expertise in place to help avoid future data breaches.

Much of this work, such as re-establishing customer confidence, involves factors beyond technology. Breached organizations are making progress on their ability to recover, with 42% reporting complete recovery this year, up from 35% last year. For most organizations, however, the hard work of recovery can take months or even years.

Figure 18.
Share of recovery from breaches

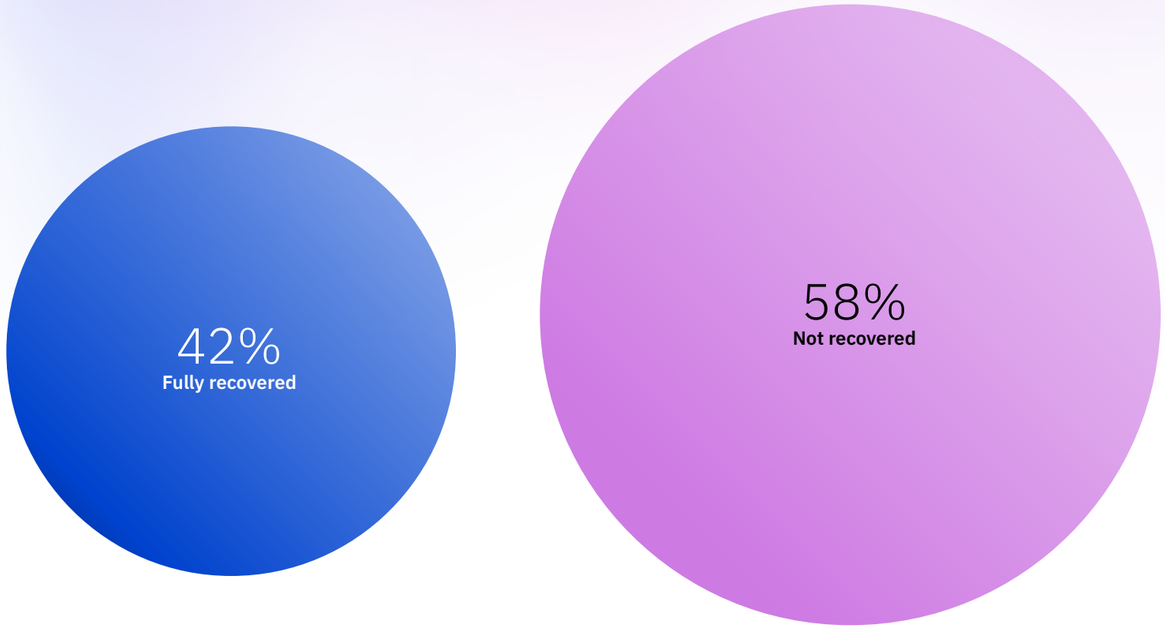

Breach recovery rates improved for a third straight year
Nearly half of organizations (42%) fully recovered from their breaches, up from 35% last year and quadruple the share from 2024 (12%). The improvement comes despite a slight uptick in slower identification and containment times. However, fewer than one in 20 organizations recovered from a data breach quickly—in less than 50 days. The biggest improvement in recovery times came at the other extreme. The share of organizations that required over 150 days to fully recover from a breach shrank to 19% from 26% last year.
See Figures 18 and 19.

Figure 19.
Share of organizations that fully recovered from a data breach

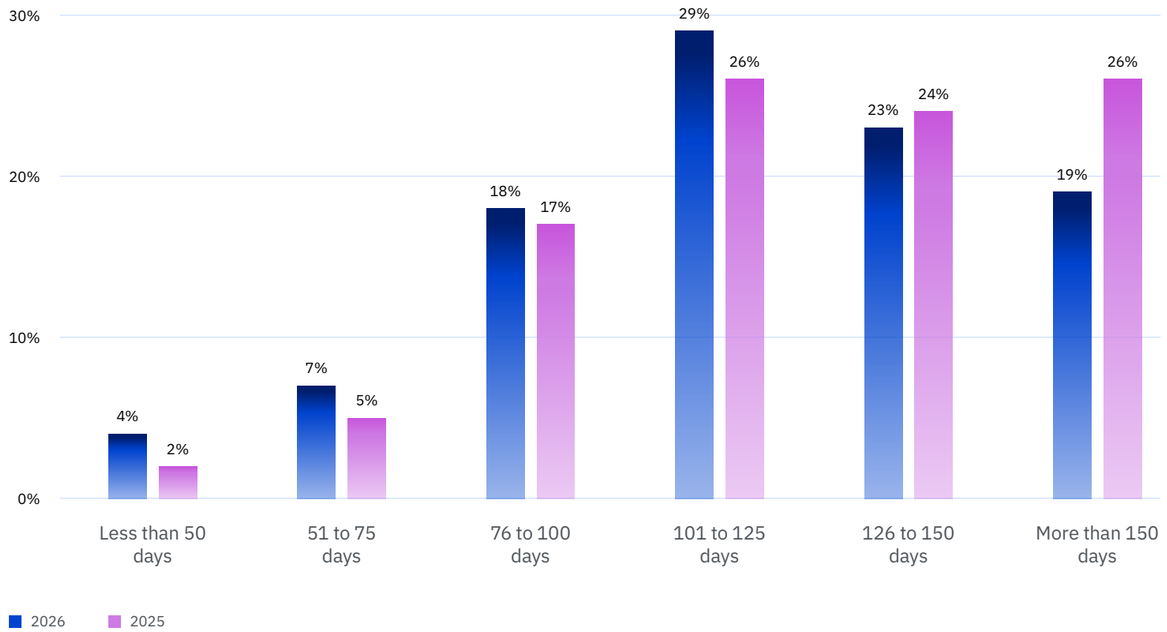

AI-driven attacks

Attackers are using AI as a critical tool in their toolkit for generating phishing emails, scanning code for vulnerabilities and automating attacks at scale. For the second year, our research looked at the scale of AI-driven attacks to understand the cost and scope and the industries most at risk.

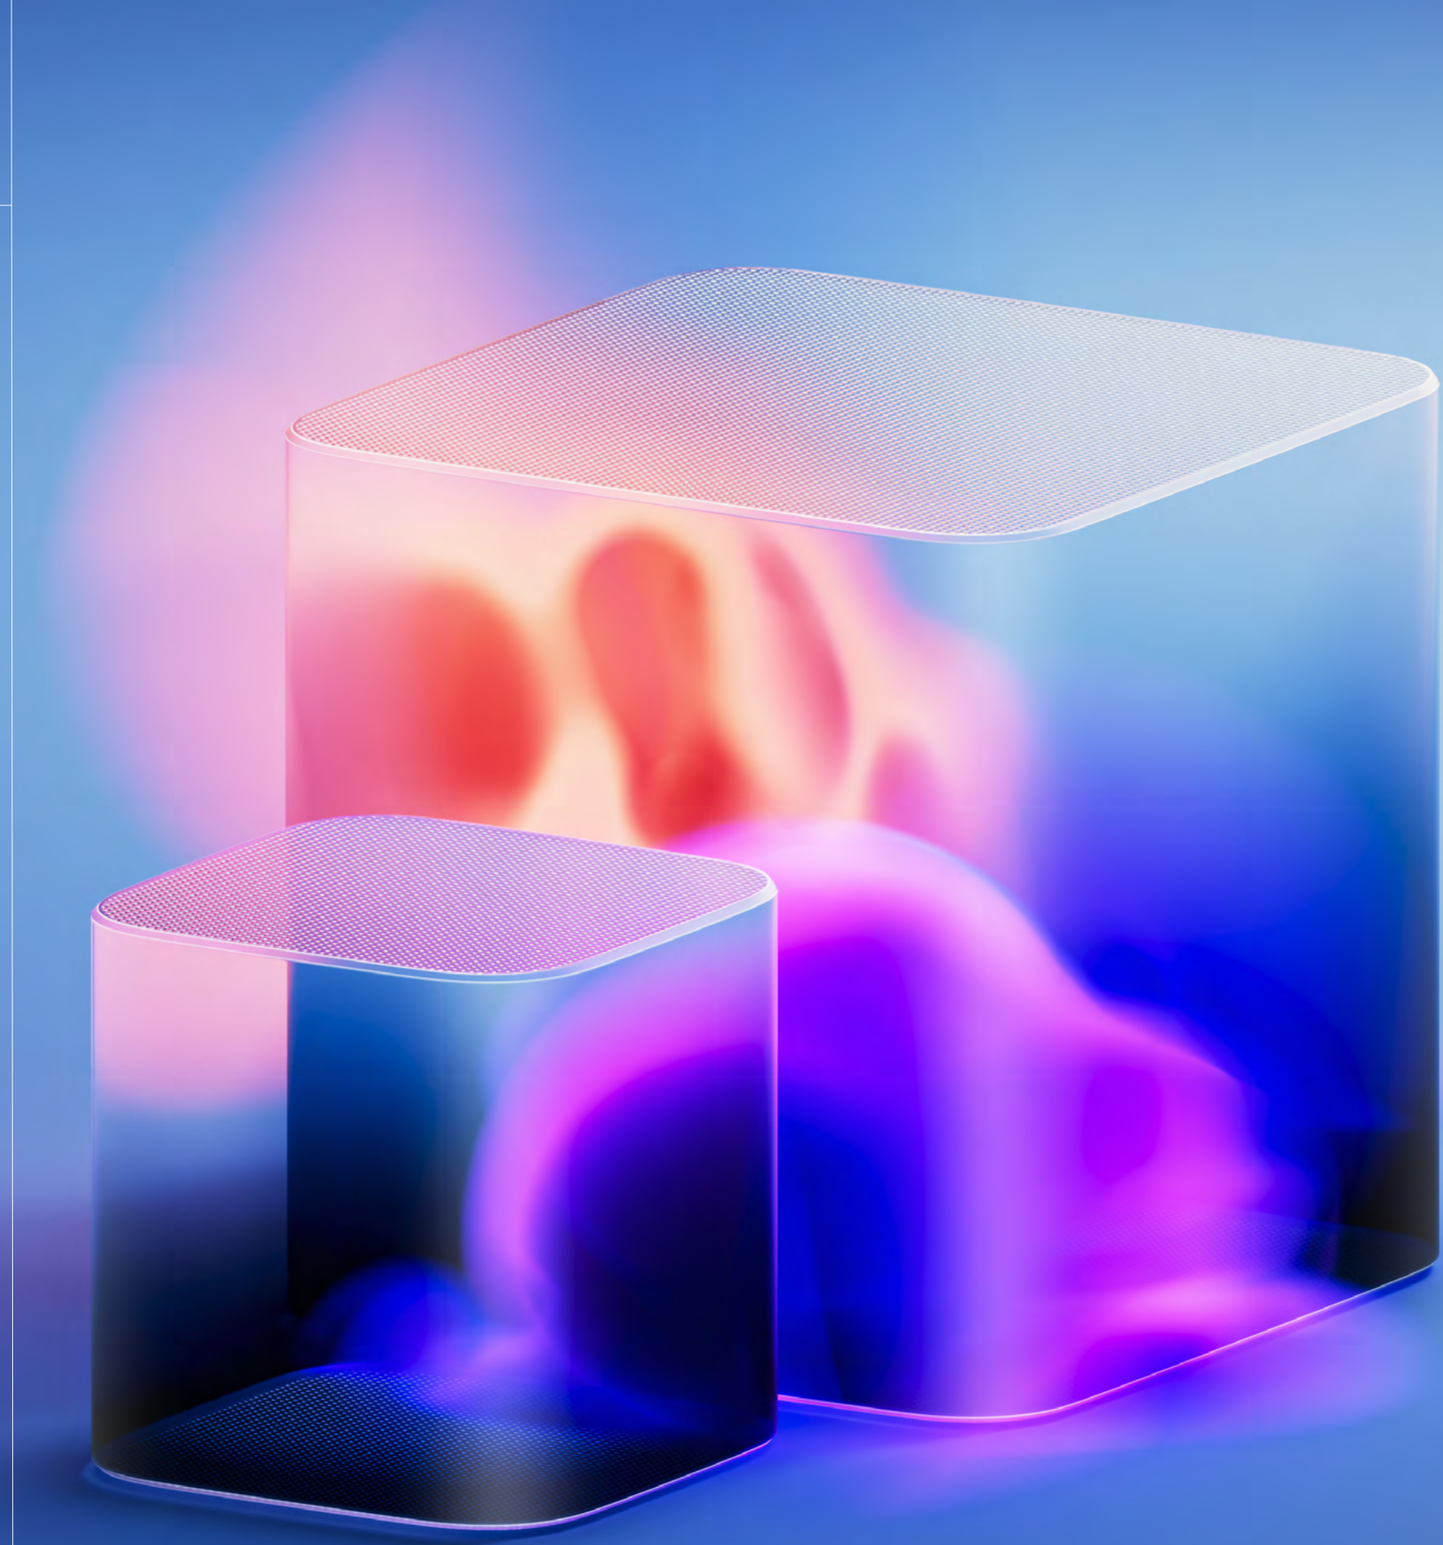

56%

Increase in AI-driven attacks

Attackers are weaponizing AI at scale
More than one in 4 organizations experienced a malicious, AI-driven attack—a 56% increase over last year. Deepfake impersonations drove the highest volume (45%). Because gen AI makes social engineering attacks cheap to create and hard to detect, attackers are increasingly using them to trick users and bypass identity controls. AI-generated malware was responsible for 19% of AI-generated attacks. IBM researchers have noticed this type of malware becoming more common.
See Figures 20 and 21.

AI-driven attacks led to higher breach costs
AI-driven attacks led to significant financial impacts, adding USD 1 million to the average malicious breach cost. As attackers adopt more advanced frontier AI model capabilities, these attacks will likely reshape breach economics by accelerating attack speed, scaling targeting and increasing downstream business impact. That speed can compress timelines from weeks to days and push the cost of delayed response even higher.
See Figure 22.

Figure 20.
Share of AI-driven attacks

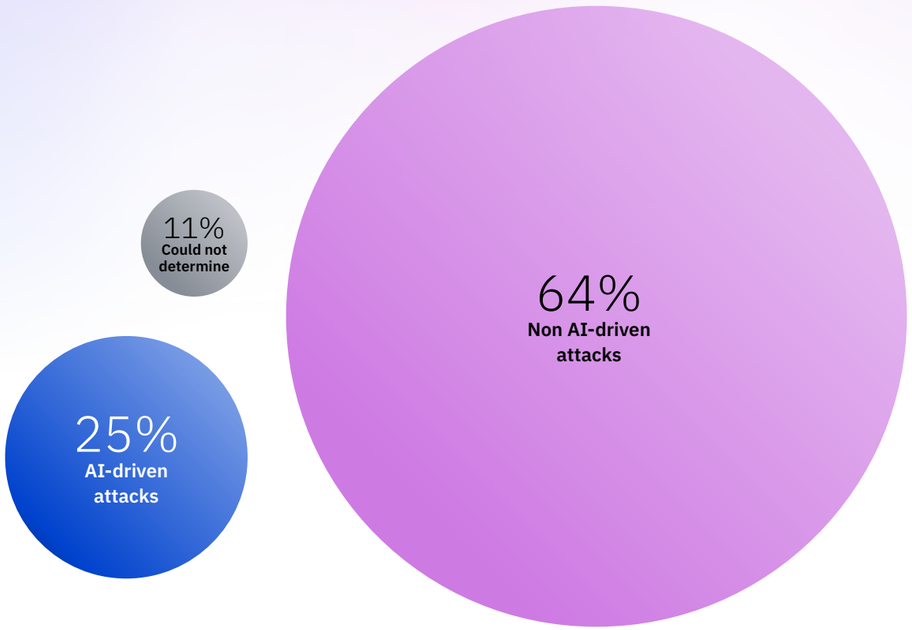

1M

Additional cost for AI-driven attacks, above the global average for malicious data breach costs, measured in USD

AI-driven attacks concentrated on critical infrastructure
Most AI-driven attacks targeted critical infrastructure sectors (a combined 62%), with financial services and energy organizations experiencing the highest concentration, raising the risk of broader systemic disruption. The concentration of AI-driven attacks on these sectors, or on a single organization in one of these sectors, is concerning. A successful attack can cascade into broader impacts across consumer finances, economic systems and power grids.

Figure 21.
Share of malicious AI attacks by type

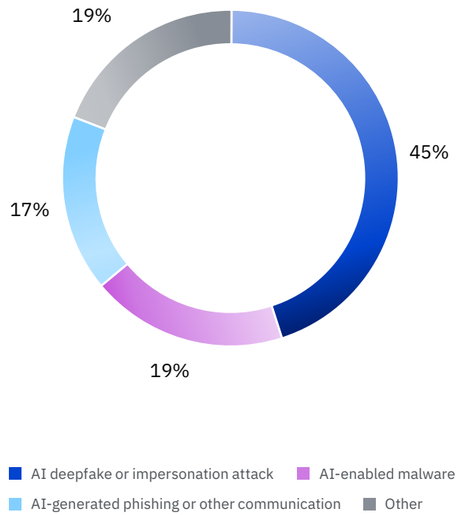

Figure 22.
Measured in USD millions

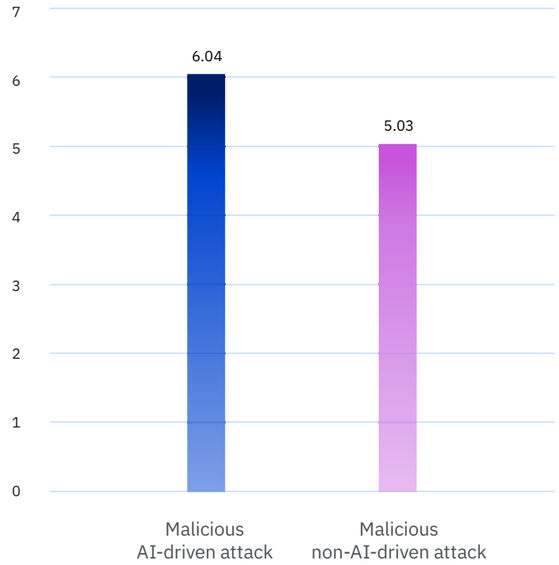

Ransomware attacks

Ransomware can have devastating downstream effects, negatively impacting an organization's operations and disrupting people's ability to access personal finances, healthcare and energy.

Thanks in part to ransomware as a service, these attacks continue to increase. But our research found attackers are shifting their tactics. In addition to encrypting data and demanding a ransom to unlock it, they're weaponizing brand reputation, increasingly resorting to publicly shaming victims to extort them.

Ransomware attacks continue to rise
Among breached organizations this year, 39% reported their systems were hit by ransomware. This finding continues a four-year upward trend, from 24% in 2023, and represents a 62.5% rise during that time.

Figure 23.
Percentage of ransomware attacks among breached organizations

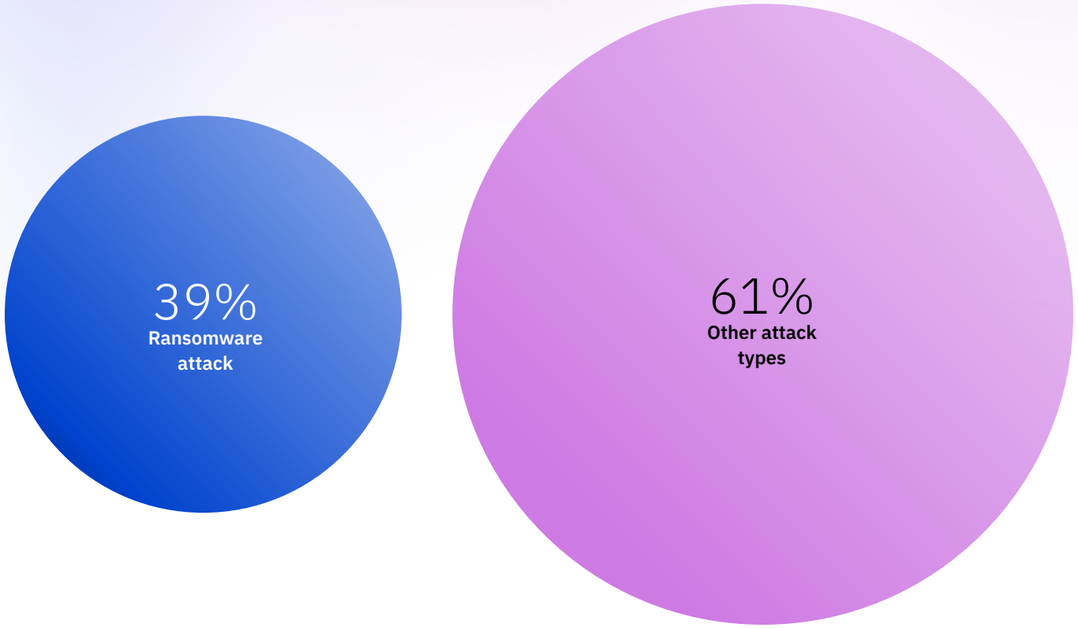

Ransomware attacks targeted PII and weaponized reputation
High-value data, such as IP and employee PII records, is still being targeted. However, ransomware attackers have broadened their scope to include internal emails and communications such as Slack messages in 19% of attacks. Traditionally, ransomware attackers relied heavily on email and focused on encrypting data to extort ransom payments. The implicit threat was data theft and disrupted operations. Now attackers use more points of pressure. Attackers are also threatening brand reputation, with 41% of attacks this year including threats of public shaming and leaking data to the media. See Figures 23 and 24.

Figure 24.
How attackers weaponize ransomware; more than one response permitted

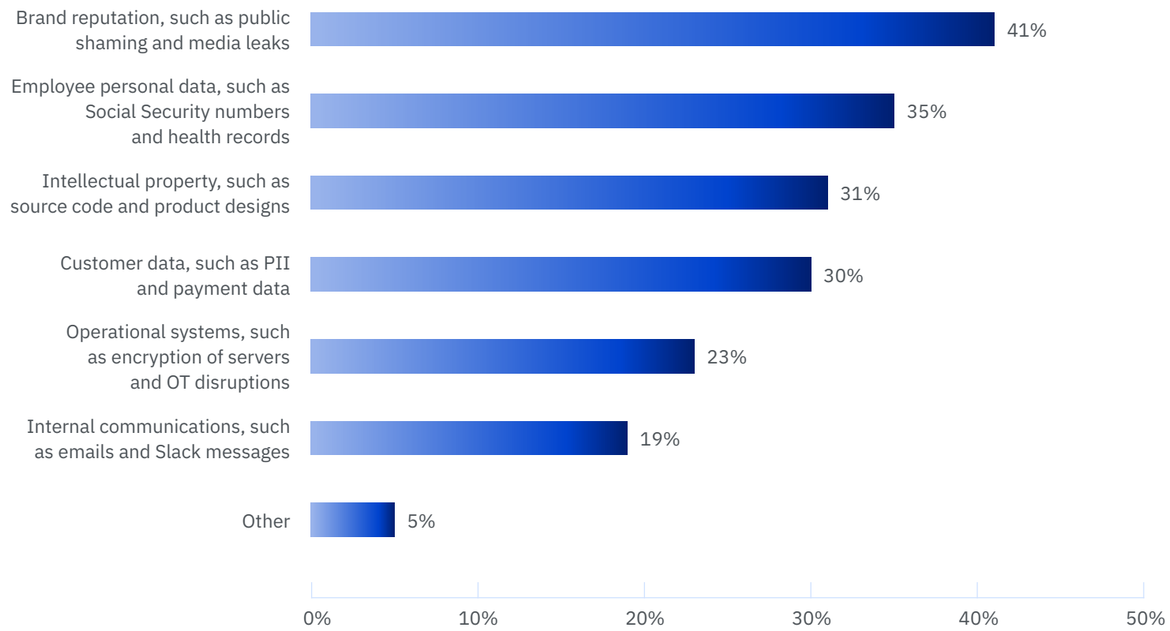

Breaches involving AI models and applications

Organizations continue to prioritize innovation over security for AI models and applications, which can leave them vulnerable to AI-related breaches. The share of those breaches grew substantially this year. These breaches compromised data, disrupted operations, harmed reputations and led to direct financial loss. The most common causes weren't model failures themselves, but weaknesses in surrounding systems.

Measured in USD

6.07M

Average cost of a breach that involved an AI model inversion incident

5.89M

Average cost of a breach that involved an AI model prompt injection incident

5.25M

Average cost of a breach that involved cloud misconfigurations affecting AI workloads

4.94M

Average cost of a breach that involved a malicious model

4.72M

Average cost of a breach that involved model evasion

Security incidents involving AI increased
AI-related breaches grew to 21% this year from 13% last year. That's a 61% increase, and it suggests attackers continue to probe organizations' AI systems and surrounding environments for weaknesses. Some organizations are failing to notice this trend. Among those that experienced an AI-related breach, 92% lacked proper AI access controls. Those controls include role-based access and multifactor authentication. The cost of these breaches was also higher than those that didn't target AI or where the organizations were unsure if AI was targeted.
See Figures 25 and 26.

Figure 25.
Share of organizations that experienced a security incident involving an AI model or application

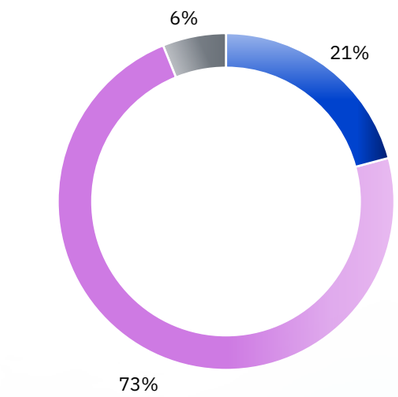

Figure 26.
Average cost for a security incident involving an AI model or application; measured in USD millions

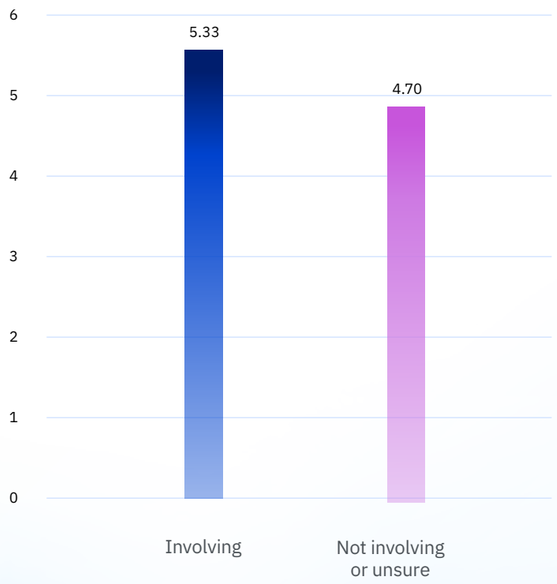

Model inversion breaches were costliest
AI-related breaches took a variety of forms, including data poisoning (corrupting training data), model inversion (extracting sensitive data from a model) and prompt injection (maliciously altering prompts to manipulate the model's output). Average costs varied by incident type. For example, a breach involving a model inversion (USD 6.07 million) was 18% higher than the global average. Prompt injections were nearly as costly (USD 5.89 million), possibly because corrupting the AI model also corrupted all the business processes that relied on it.
See Figure 27.

Figure 27.
AI-related breaches by incident type; measured in USD millions

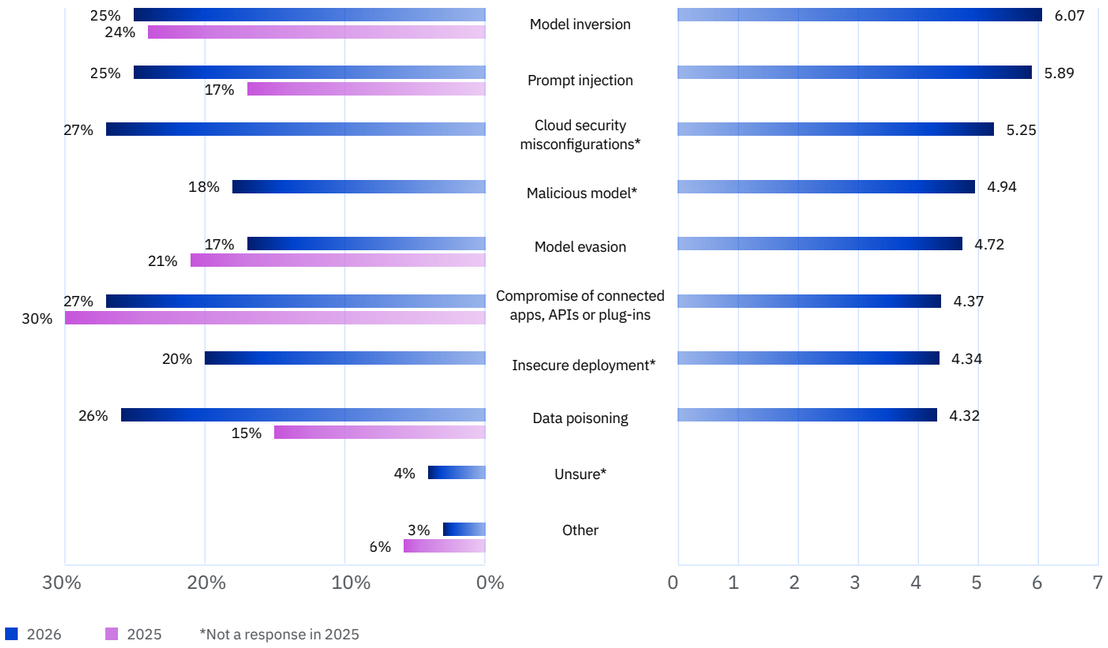

Financial loss was the top impact of AI-related breaches
AI models and applications have been integrated into many daily workflows and are responsible for helping drive business operations. That integration may explain why financial loss is one of the top impacts for AI-related breaches. More than half of organizations reported financial losses from these breaches, while 44% reported operational disruptions and 26% suffered reputational damage. See Figure 28.

Figure 28.
Impacts of AI-related breaches; more than one response permitted

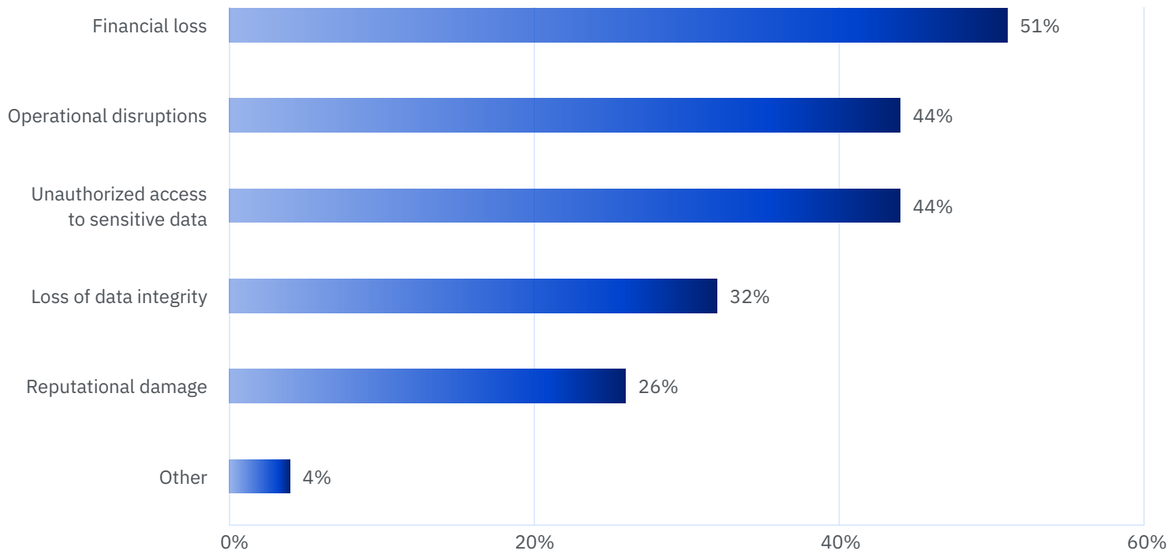

AI-related breaches showed no difference by deployment type, but varied by cost
AI-related breaches occurred at similar rates across deployment types, averaging 26% for open-source, and third-party vendor deployed as SaaS, and 25% for third-party vendor delivered on-premises. Breach costs by sources varied. Breaches involving in-house models averaged USD 4.98 million, and open-source models averaged USD 5.63 million. See Figure 29.

Figure 29.
Source of the AI model or application attacked; measured in USD millions

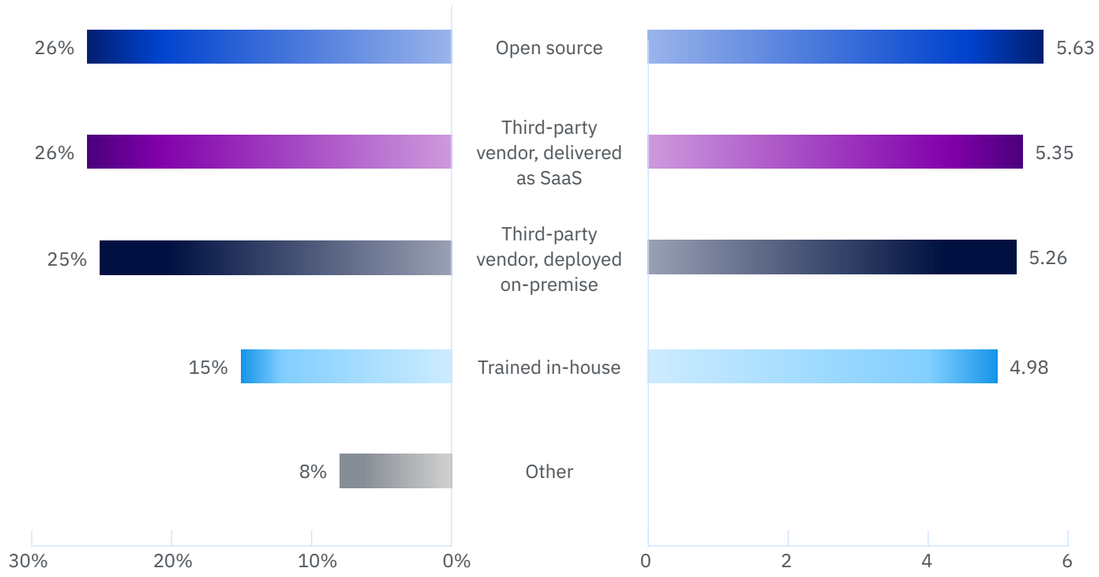

Incidents involving shadow AI more than doubled
Security incidents involving an organization’s shadow AI—where workers use unapproved AI—more than doubled to 43% this year from 20% last year. These incidents led to higher average breach costs this year (USD 5.39 million) than last year (USD 4.63 million). They also had serious consequences, including data loss or compromise (49%) and disrupted operations (42%). In this year’s research, we also looked at regulatory fines related to shadow AI breaches. In about one in five incidents, organizations reported paying a fine. See Figure 30.

Figure 30.
Impacts of a shadow AI security incident; more than one response permitted

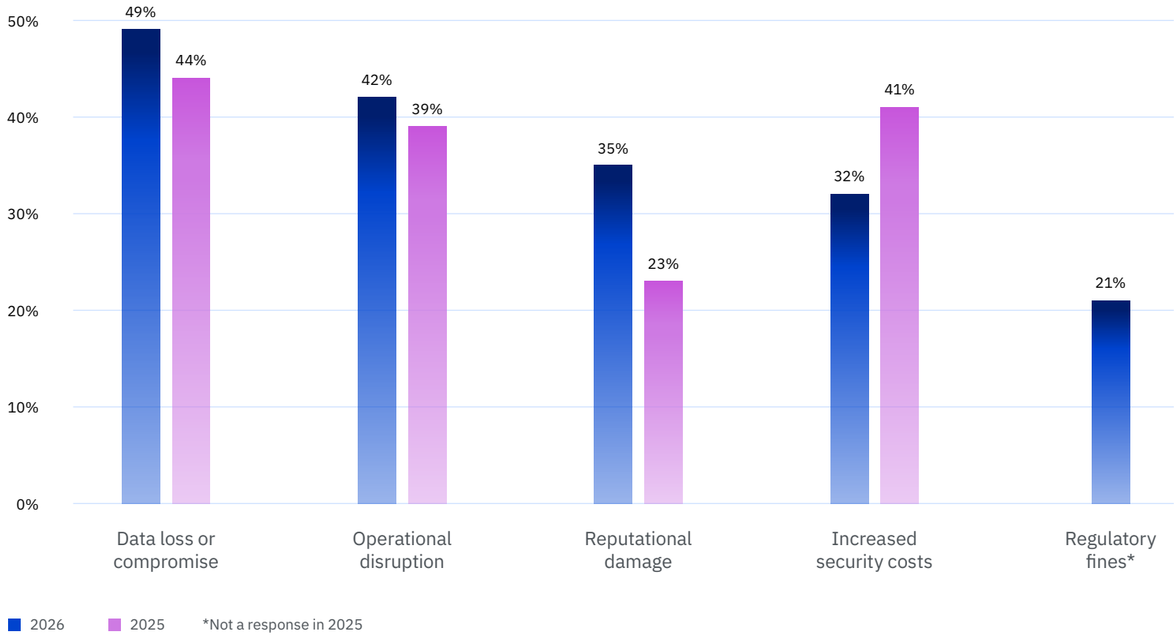

AI oversight

As we noted last year, AI adoption is outpacing oversight. That trend has only increased. More organizations this year lacked governance policies to mitigate or manage the risk to AI. Among those that had policies, only about a third had strict approvals for deploying AI. While governance and security for AI must work together, even fewer organizations reported coordination among these critical functions. Persistent oversight gaps leave organizations open to security, operational, compliance and reputational risks.

Governance for AI is being ignored
AI governance brings discipline, control and oversight to AI operations. It includes everything from access controls for models, agents and applications to monitoring AI activity for data leaks. Despite its importance, most breached organizations this year lacked AI governance to manage AI or detect shadow AI (68% vs 63% last year). On the other hand, more organizations expressed growing concern for AI governance, with a third (33%) reporting they were developing governance policies, up from 22% last year. See Figure 31.

Figure 31.
Organizations with policies in place to manage AI and prevent shadow AI

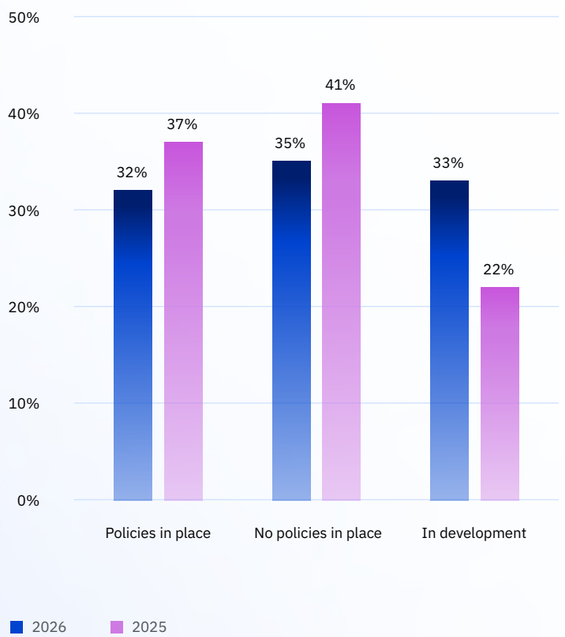

AI approval processes top governance policy types
When organizations enforced AI governance, they used a mix of policies and technologies. The most popular governance controls required IT approval for AI deployments (38%), but that critical oversight was down from 45% last year, when we first researched the issue. Following closely was deploying AI governance frameworks or technologies (33% each). One organizational challenge we studied for the first time this year involved coordinating among governance and security teams. Only 19% of organizations reported they coordinated these efforts. That lack of collaboration could lead to blind spots, policy conflicts and slower response times to security incidents. See Figure 32.

Figure 32.
Share of organizations and policy types they use to govern AI

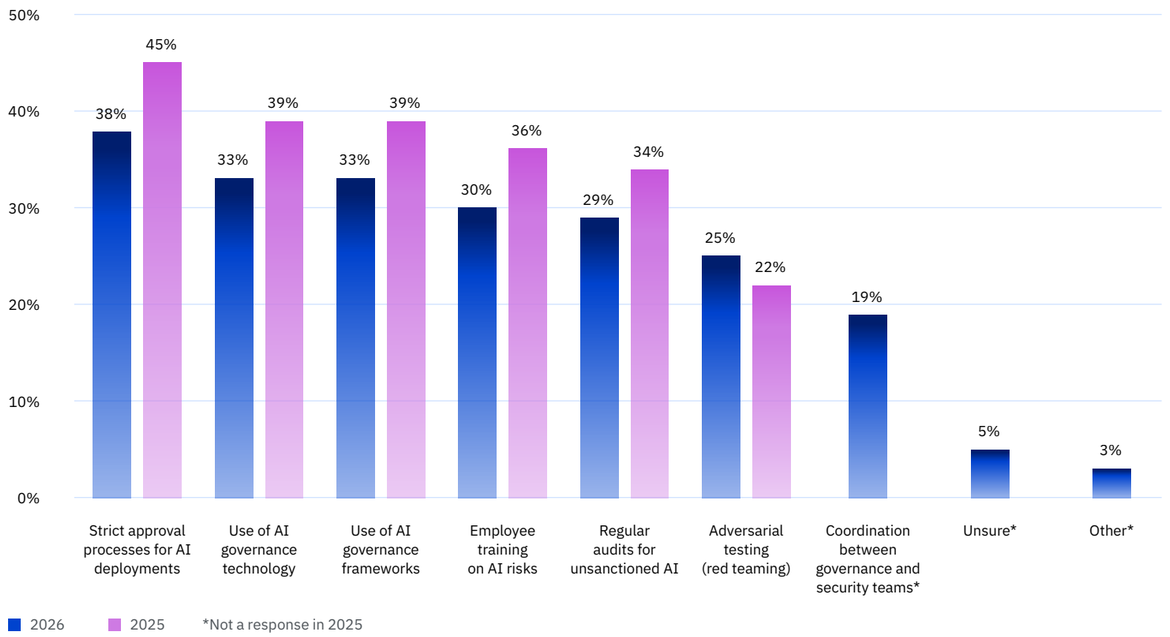

Factors that increase or decrease breach costs

When analyzing breach costs, it's important security leaders understand which technologies or events tend to lower or raise those costs. Generally, using AI and automation lower costs, but this year we wanted to take a closer look at specific technologies and use cases that affected costs. Our analysis examined 30 contributing factors and the impact of each in isolation against the global average. Also included are the top three factors found to amplify or mitigate the average data breach cost.

Key factors that reduced costs
DevSecOps approach—often a top cost reducer—again showed up as the number one factor to reduce data breach costs. Implementing an IAM solution and using an MSSP rounded out the top three factors. MSSPs delivered the fastest breach identification and containment timelines, so their appearance as a top factor in mitigating costs isn't surprising. All three of these security approaches center around and strengthen insight, intelligence and coordination. See Figure 33.

Key factors that increased costs
Supply chain data breaches, in which business partners were compromised and became an attack vector, were the most expensive factor in increasing breach costs, adding USD 227,250 on average. This sort of data breach is difficult to detect and contain, since it involves another organization's IT infrastructure and security practices. Security system complexity came next, likely because complexity makes all security operations more difficult and time-consuming. Noncompliance with regulations was third, possibly reflecting the costs of regulatory fines following a data breach. See Figure 33.

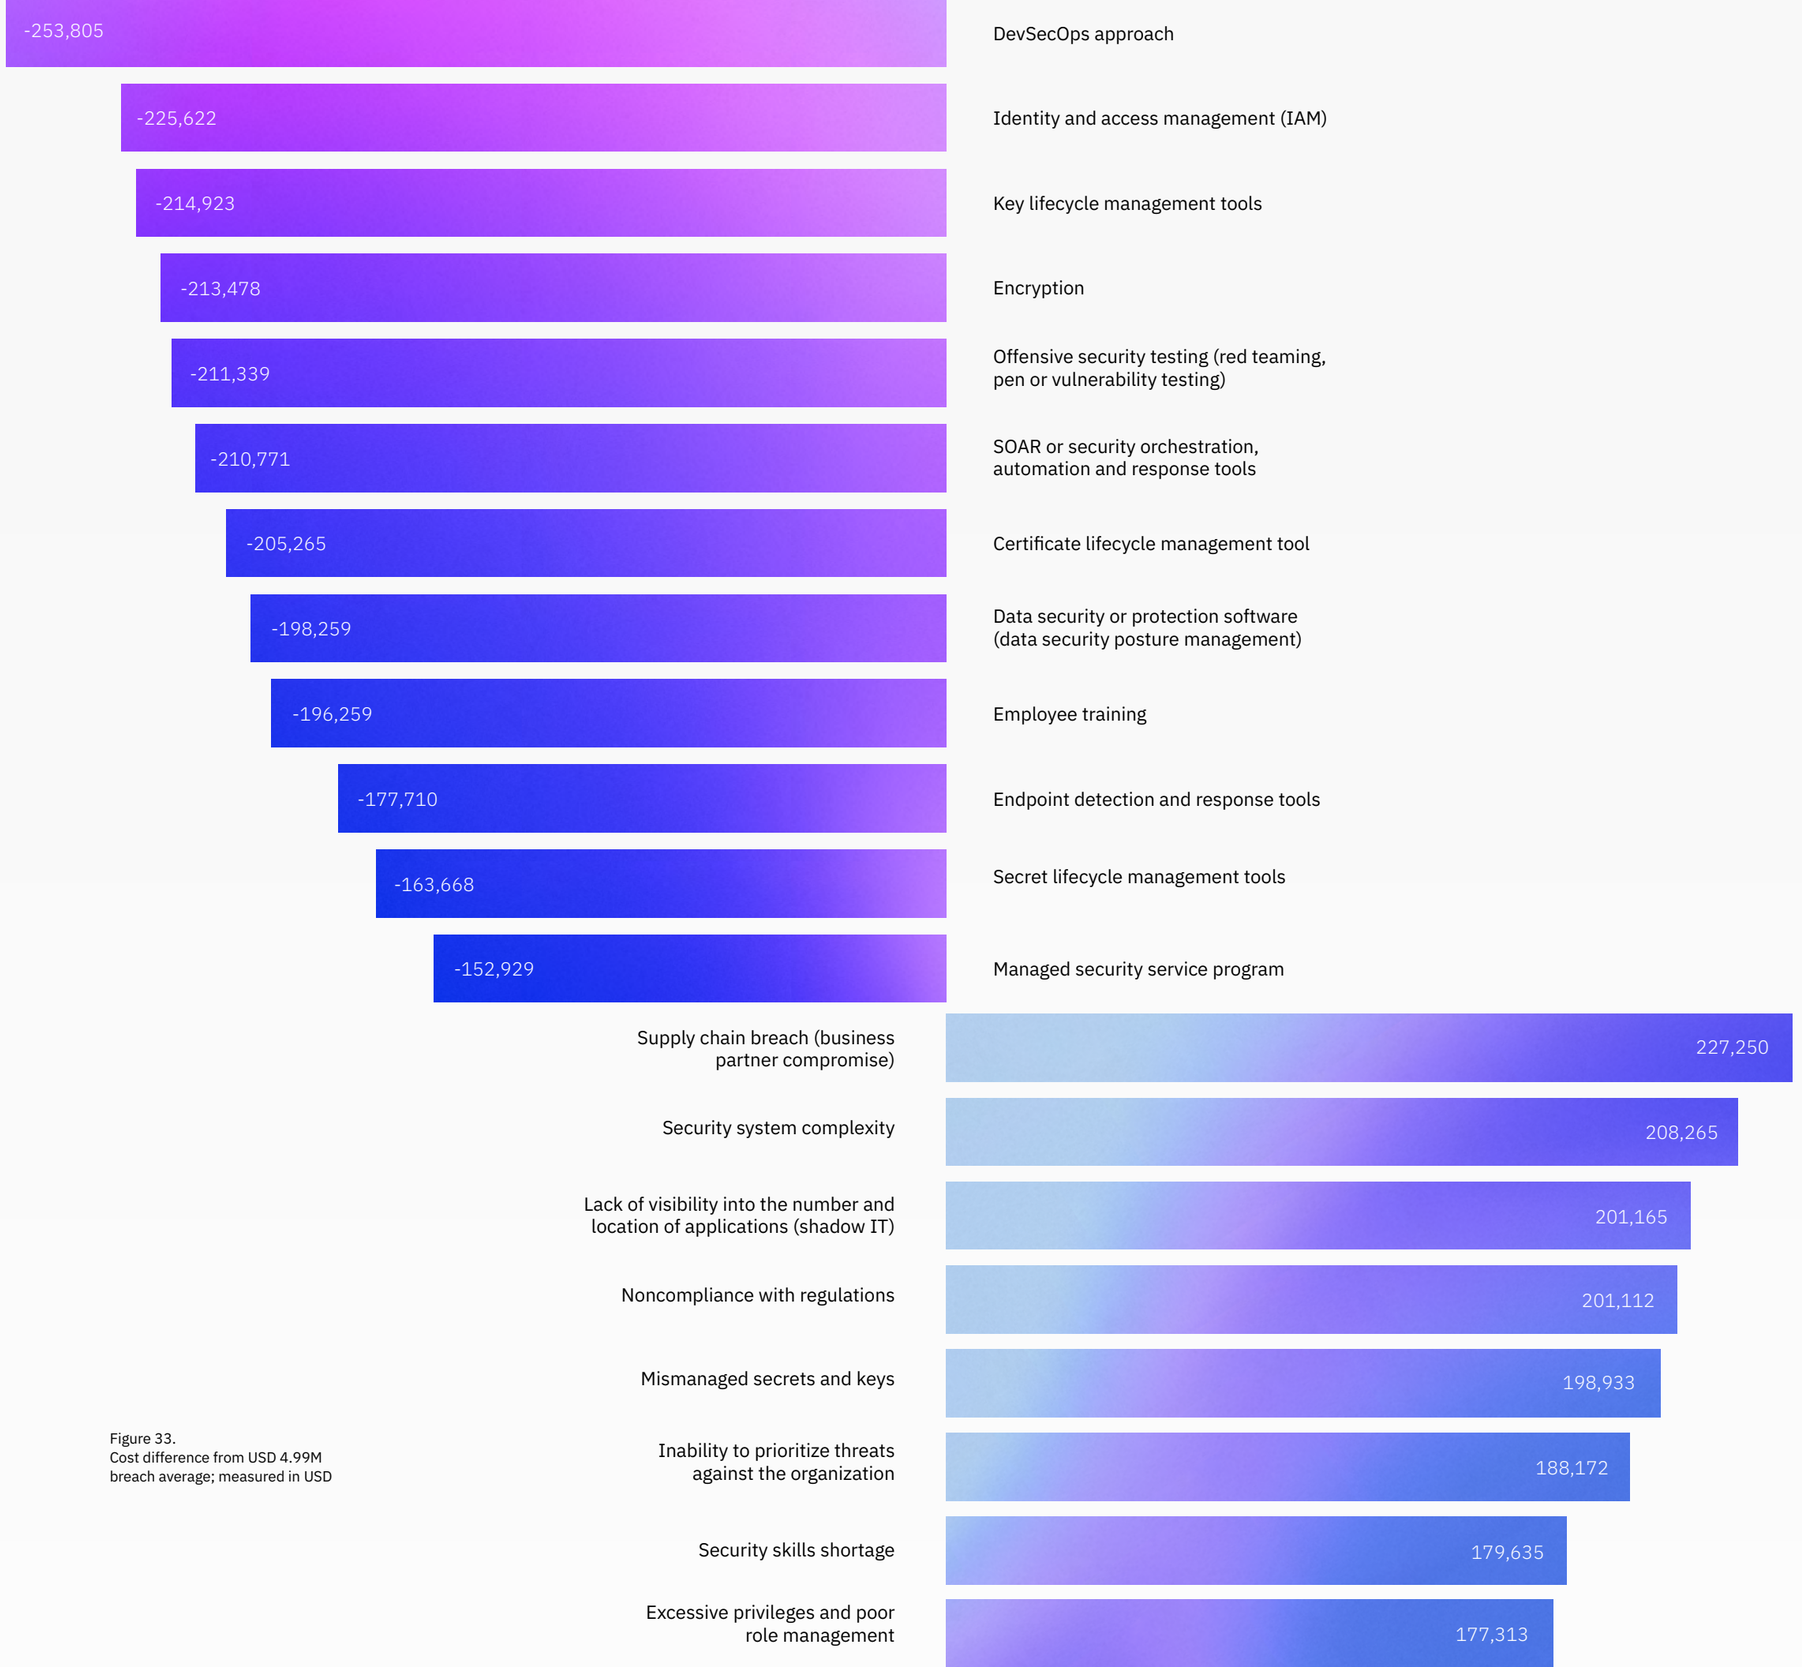

Figure 33.
Cost difference from USD 4.99M breach average; measured in USD

Security AI and automation

As frontier AI models put unprecedented speed in the hands of attackers, it's critical that defender teams use AI and automation to accelerate and scale their operations. Over the past few years, this report has noted the vital role AI is playing in security, quickening response times and lowering average breach costs. These tools can act as skills multipliers for defenders, enabling them to oversee more systems and react faster to threats. While these gains are clear, adoption rates remain static and uneven across the security functions.

Extensive use of AI and automation rose slightly
Organizations that used security AI and automation extensively remained a minority but grew this year. They accounted for 36% of the organizations we studied, up from 32% last year. However, most organizations (64%) reported limited or no use of the technologies in security operations. See Figure 34.

Extensive use of AI and automation reduced breach costs
Security AI and automation continue to function as cost reducers in breaches. Organizations that used these technologies extensively had an average breach cost of USD 4 million, while those that didn't use these technologies had a much higher average breach cost of USD 5.93 million. These figures represent a USD 1.93 million cost savings for high adopters, a nearly USD 400,000 improvement over last year. For limited-use adopters, the savings were lower but still substantial, at USD 880,000. See Figure 34.

Figure 34.
Percentage of organizations per usage level; measured in USD millions

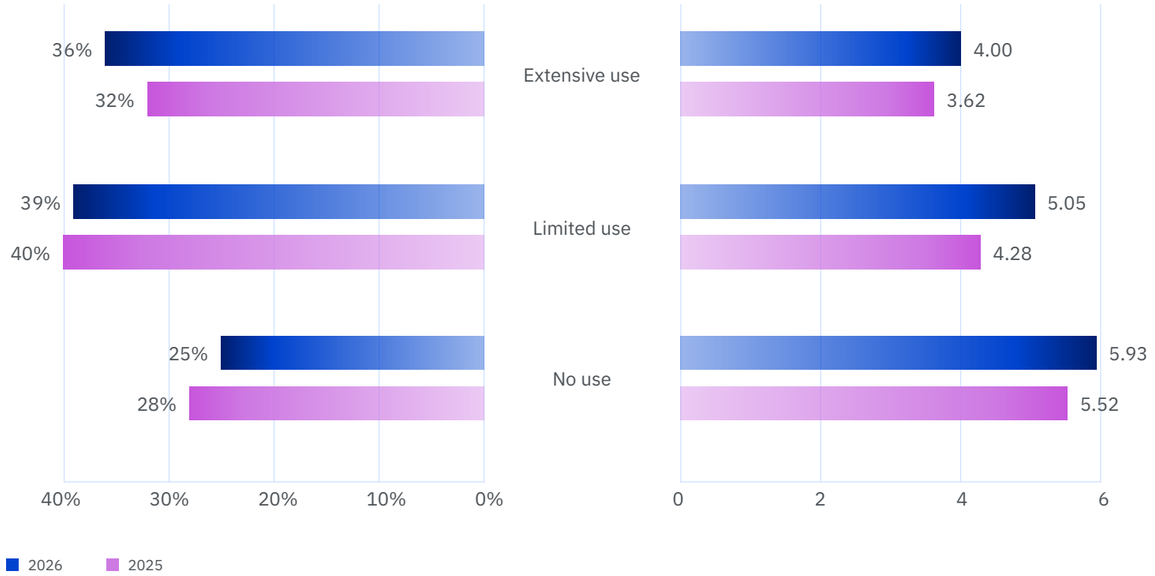

Extensive use of AI and automation shortened response times
Organizations that used AI and automation extensively shortened identification and response times to 215 days. That's 65 days faster than those not using these technologies at all (280 days). Not surprisingly, those slower times correlated to higher breach costs. See Figure 35.

Figure 35.
Time to identify and contain a breach with and without AI and automation; measured in days

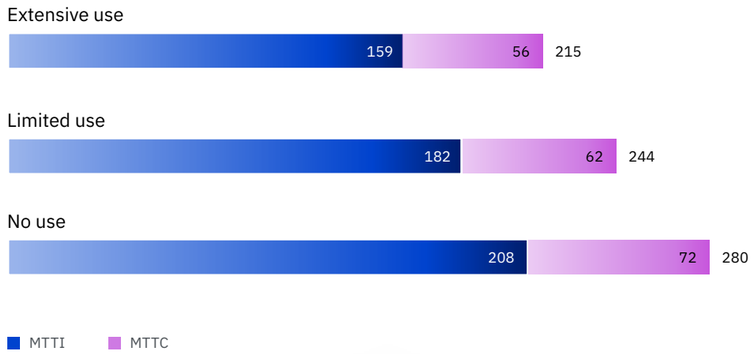

AI and automation remain uneven across SOC operations
When using AI and automation in security, organizations favored one area: threat detection and response. Among breached organizations, 77% made either limited or extensive use of these technologies to investigate threats. Conversely, their use in threat prevention was lagging, with a combined 69% of use in that area. That means more than a third of organizations weren't using these technologies in their front-line security efforts, identifying and patching vulnerabilities before they can be exploited. This work is vitally important given the ability of new frontier AI models to discover vulnerabilities faster than traditional patching cycles can address them. See Figure 36.

Figure 36.
Percentage of organizations per usage level

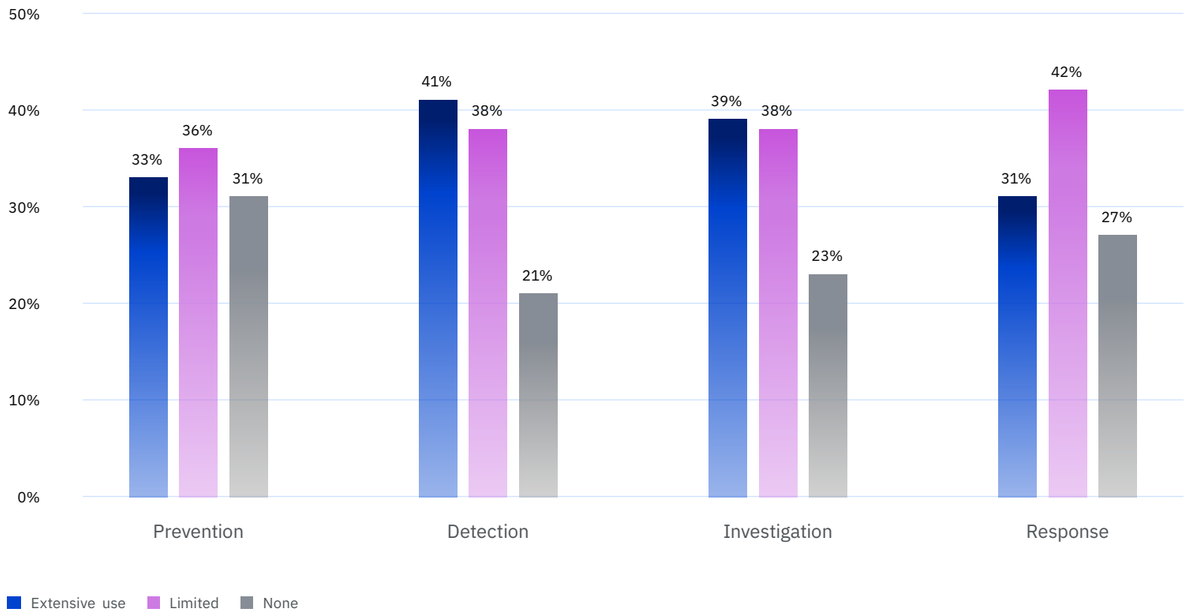

Agentic AI in security

For the first time, our research investigated AI agent deployment in the SOC and found a pattern similar to the adoption of AI and automation technologies.

A minority of security teams that used agents prioritized their use in detection and response over prevention. As attackers use AI to shorten exploit windows, this gap leaves known exposures unresolved longer, driving higher breach costs and widening the advantage for attackers. This growing imbalance—where attacks can be launched for thousands while breaches cost millions—is reshaping how organizations think about cyber risk.

Security teams using agentic AI adoption favored it in two categories
Among the 50% of organizations that have an SOC and deployed agents in it, the majority are using those agents for threat hunting and response and containment. Only 18% are using agents in the SOC for vulnerability scans and management. This area is where security teams most need a machine-speed advantage to counter the new frontier AI model threats. See Figure 37.

Figure 37.
Percentage of organizations by AI agent adoption category

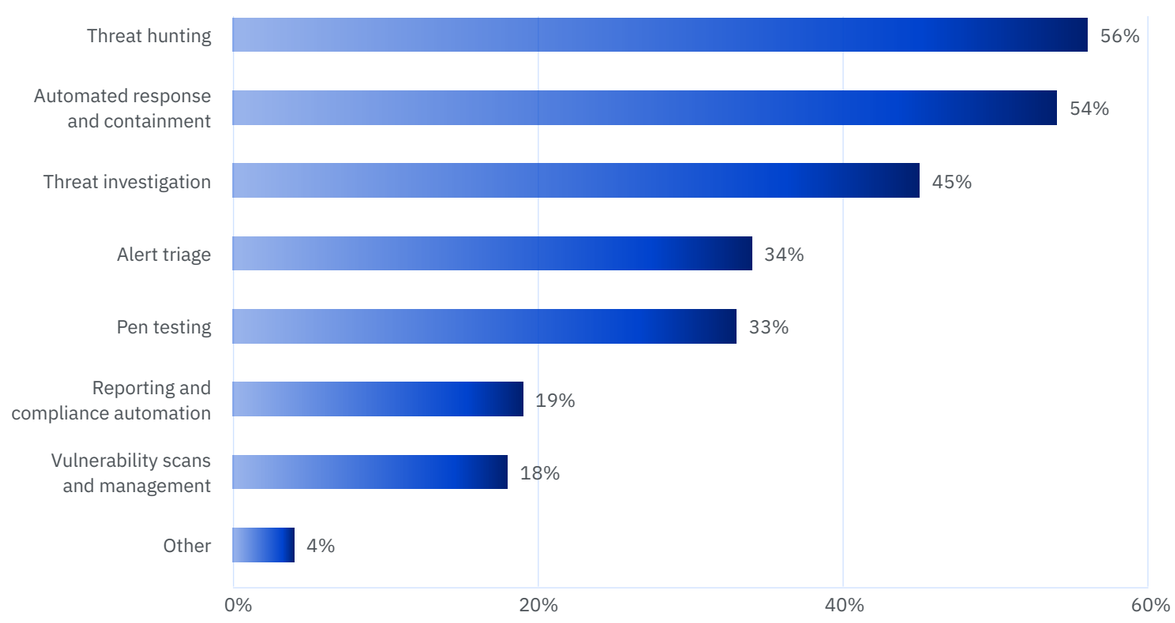

Security investments

Each year, organizations are asked if they plan to fortify their security defenses following a breach and, if so, in what areas. This year, researchers went back to those organizations and found a sharp spike in planned increase after organizations had become aware of advanced frontier AI model cyber capabilities. These models are forcing organizations to start acting on perceived future risk, rather than waiting for an incident.

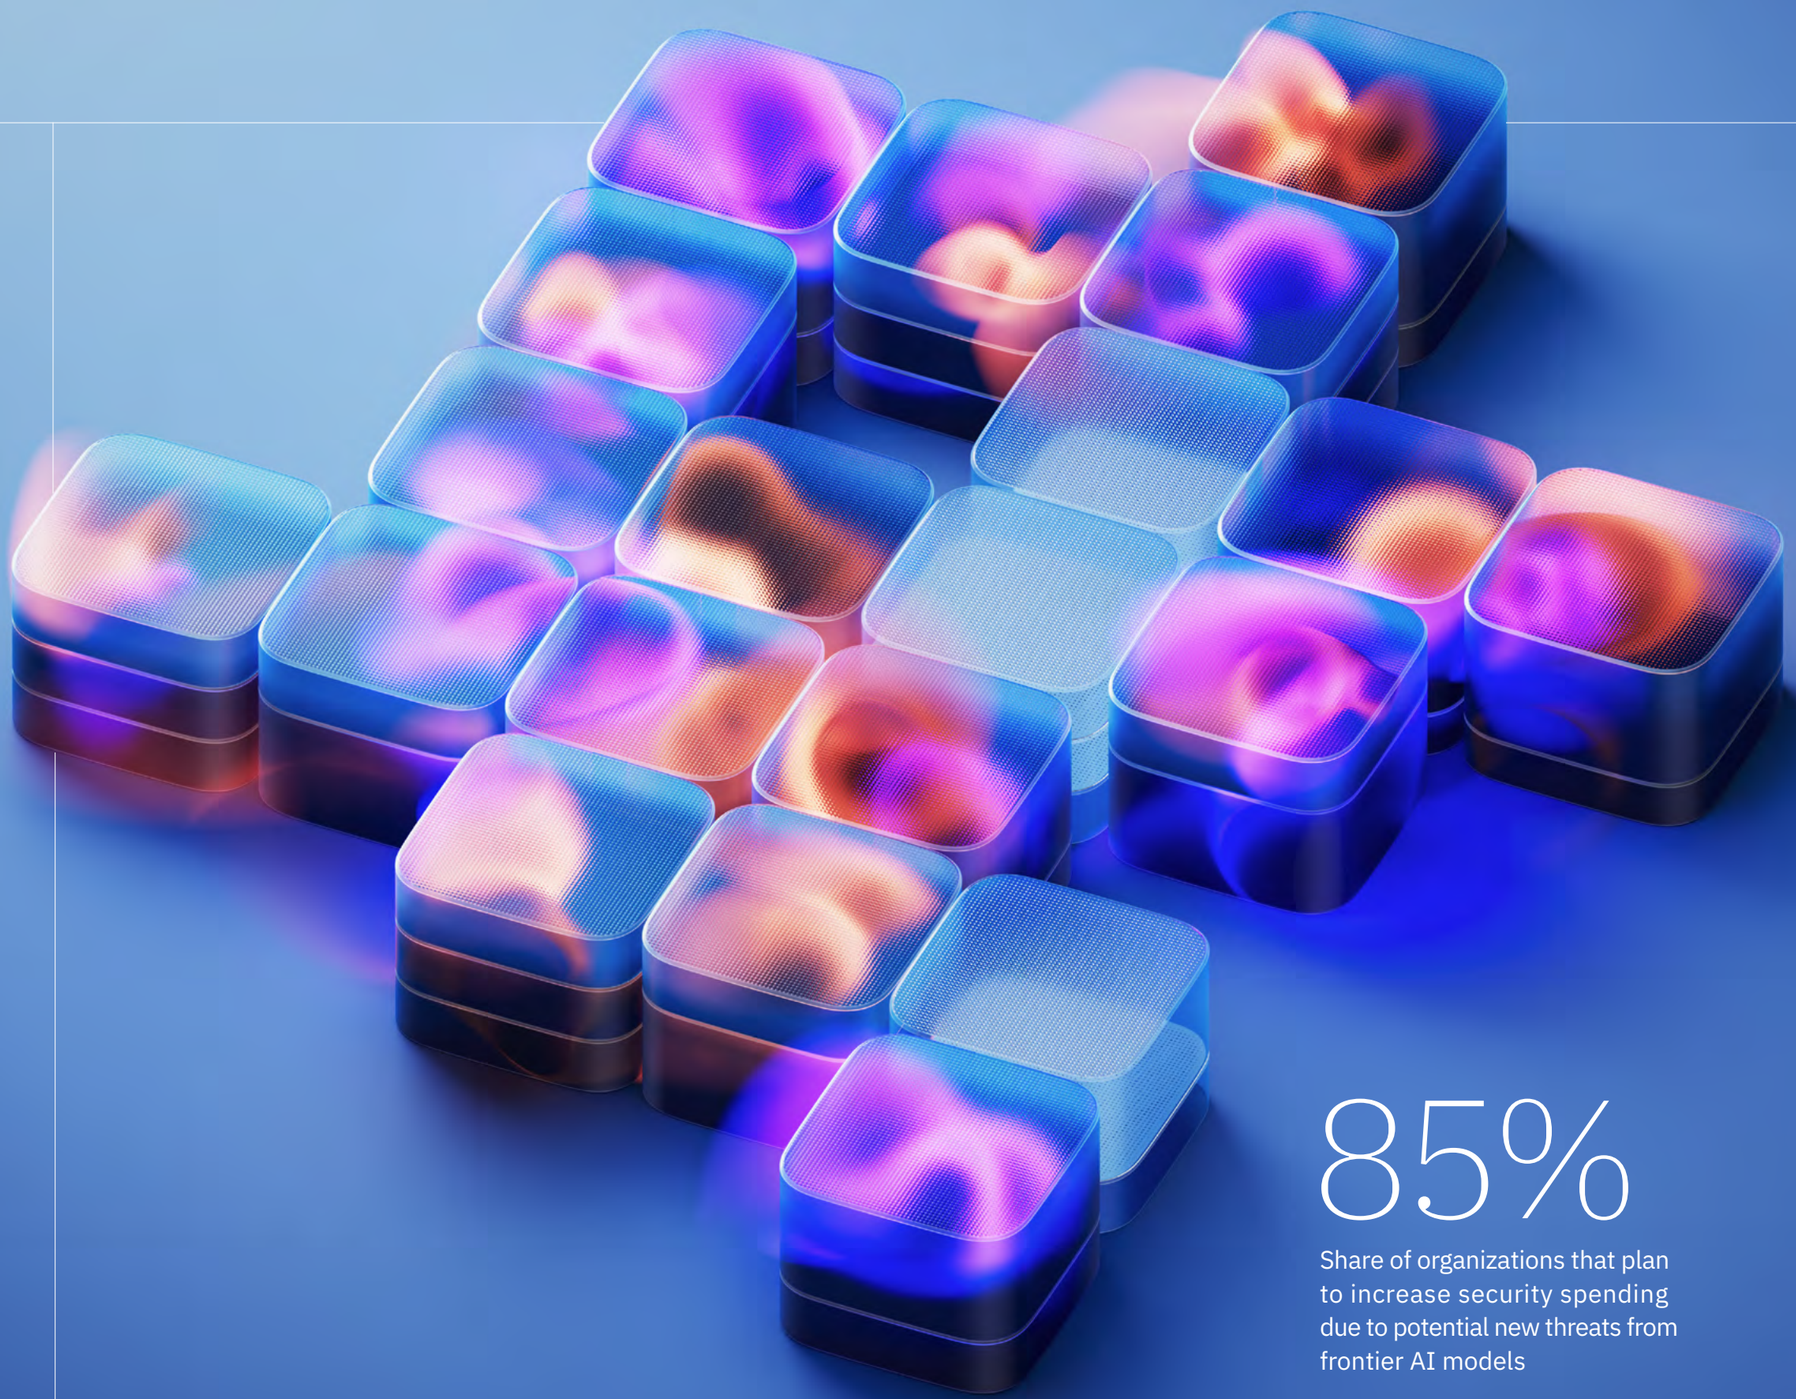

85%

Share of organizations that plan to increase security spending due to potential new threats from frontier AI models

Investment plans increased in response to frontier AI model threats

More than half of organizations (64%) said they would increase security investments following a breach. That figure is a 30% increase over last year, when investment planning was in the minority, and declined. However, in follow-on research, 85% of the breached organizations that were aware of new frontier AI model threats said they would increase their security spending because of it. This research underscores a shift toward proactive defense as attacker capabilities rise. See Figures 38 and 39.

Figure 38.
Percentage of organizations that will increase security spending due to new frontier AI model threats

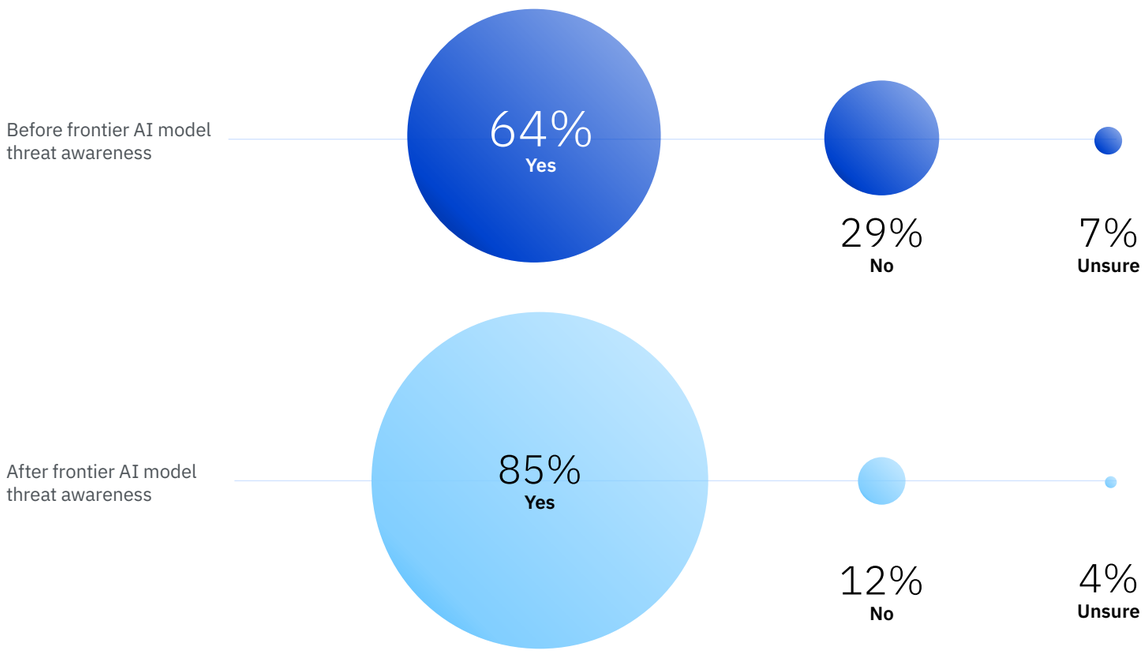

Figure 39.
Areas of increased investment before frontier model threat awareness

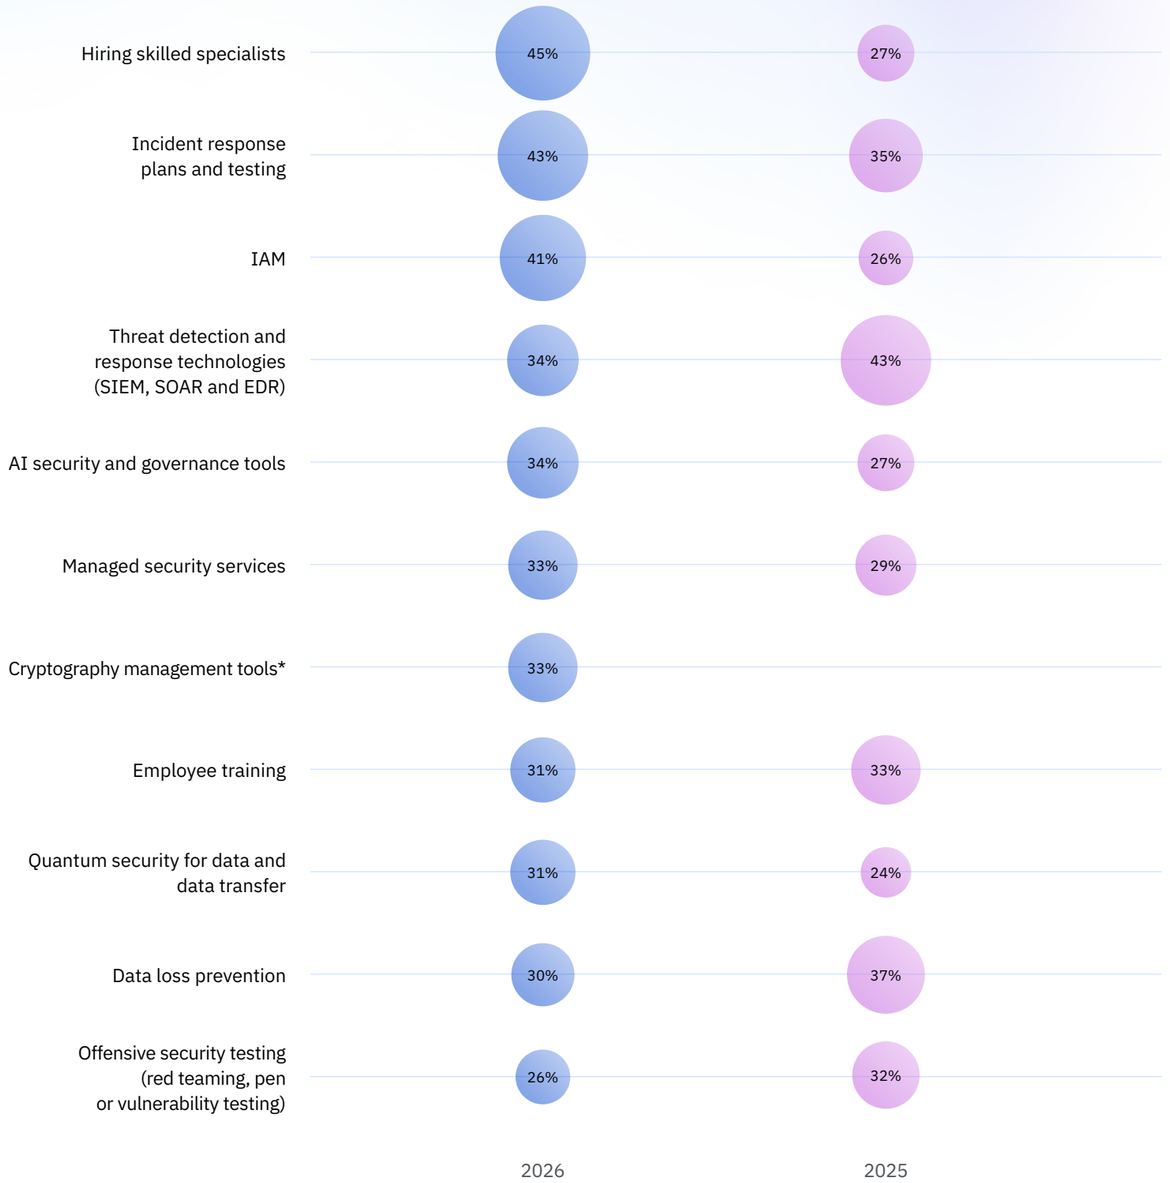

*Not a response in 2025

Agent adoption plans increased in response to frontier AI model threats

Because of threats from frontier AI models, a majority of organizations (74%) have rethought if and where they deploy agents in their SOCs. In follow-on research, most organizations said they plan to deploy agents in areas where they were underrepresented, such as in alert triage (60% versus 34%), penetration testing (52% versus 33%), and vulnerability scans and management (37% versus 18%). See Figures 40 and 41.

Figure 40.
Share of agentic AI deployment plans affected by new frontier AI model threats

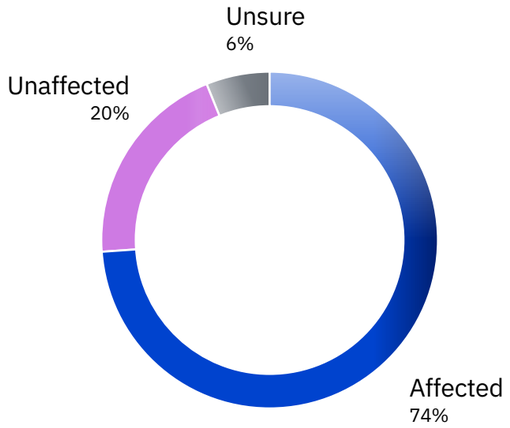

Figure 41.
Areas of agent use

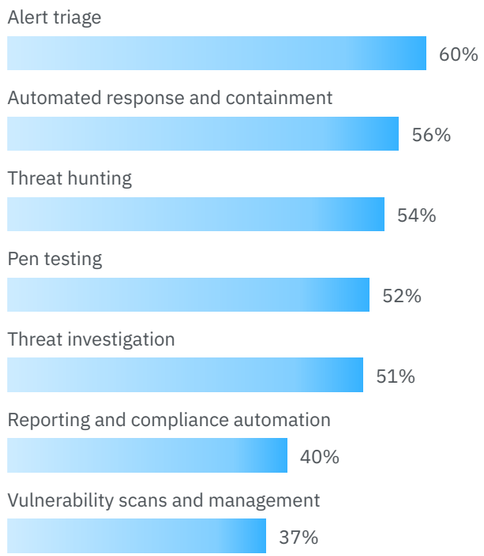

Securing non-human identities

Non-human identities (NHIs) are digital credentials assigned to machines, devices and other non-human resources.

Many organizations prefer to secure processes with NHIs rather than employee credentials. That way, the security of any process doesn't depend on the employment status and security habits of a single employee. Security teams can then track NHIs to ensure they're secure and that processes using those NHIs are behaving properly. But if NHIs themselves aren't secure, business processes that rely on them are at risk.

Most organizations don't secure NHIs in AI workflows
Less than half of organizations studied (46%) report securing NHIs in AI workflows. Among those that do, the majority (55%) use machine identity and lifecycle management. Other approaches included secrets management systems (39%), behavioral monitoring tools (36%) and role-based access controls (30%). Given the importance of securing AI workflows, taking a multilayered approach is prudent. See Figure 42.

Figure 42.
Share of NHI monitoring and control types among breached organizations; more than one response permitted

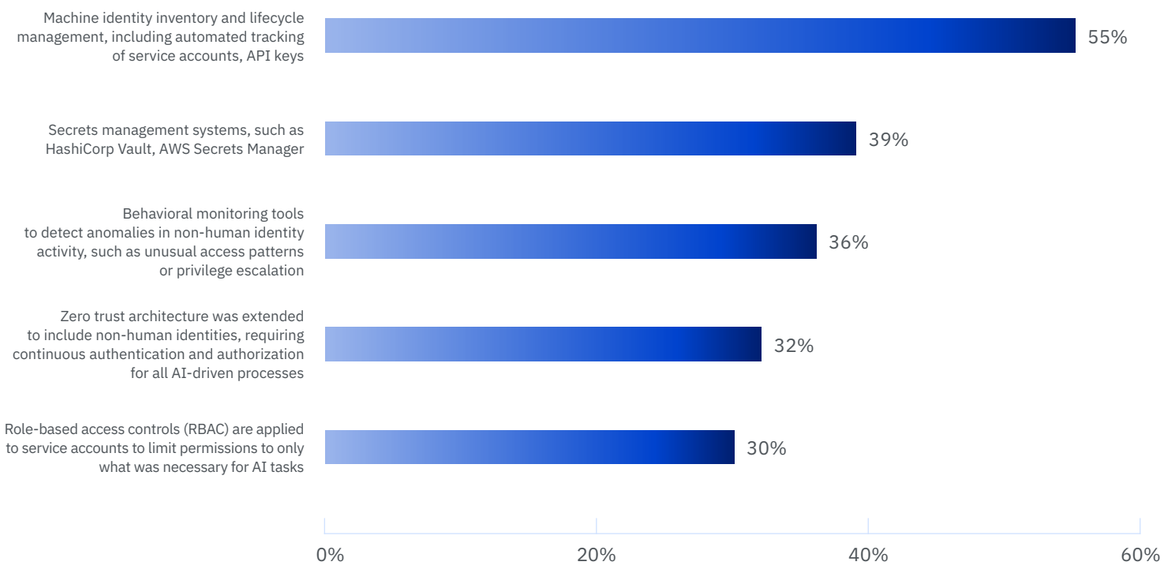

Post-quantum cryptography

As quantum computing moves closer to mainstream reality, security teams are grappling with the risks it poses, namely its ability to break long-held encryption methods, such as Rivest–Shamir–Adleman (RSA) and elliptic curve cryptography (ECC).

For the first time in our report, researchers asked breached organizations about the status of their post-quantum cryptography projects, the cryptographic algorithms designed to secure organizations against quantum computer attacks.

Most organizations don't have a post-quantum cryptography plan
Roughly a quarter (26%) of organizations reported having a post-quantum cryptography project. Most (69%) didn't and 5% were unsure if they did. Both the lack of a plan and lack of clarity can leave organizations exposed to "harvest now, decrypt later" breaches of sensitive data and company IP. See Figure 43.

Most organizations lack controls to secure cryptography
Most organizations (61%) reported that they lack controls to monitor and secure cryptography and cryptographic objects, such as keys, certificates and algorithms, across their environment. This security gap can be attributed to outdated cryptography and the expanding attack surface. Securing these assets is essential to prevent data breaches, achieve crypto-agility and mitigate future quantum computing threats. See Figure 44.

Figure 43.
Percentage of breached organizations with a post-quantum cryptography project

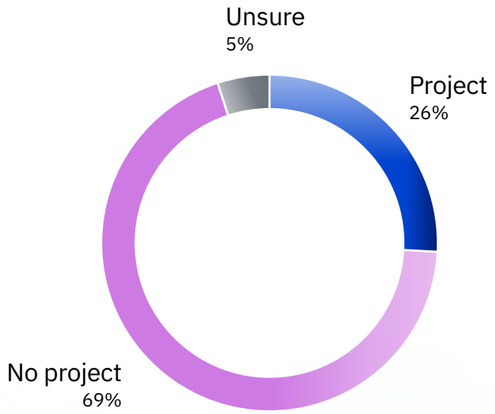

Figure 44.
Percentage of organizations that have controls to secure cryptography

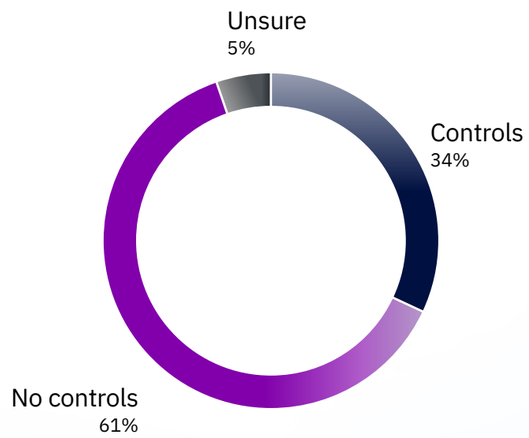

Recommendations

As frontier AI models collapse the timeline between vulnerability discovery and exploitation—propelling attacks from human to machine speed—the cost of delay will be measured in minutes, not months.

Failures in detection delays, containment gaps, lateral movement, credential abuse and decision latency are now exponentially more damaging and costly to organizations. This year's research underscores this finding.

To help prevent, mitigate and reduce the costs of a data breach—and to prepare for the age of machine-speed security—IBM experts suggest these four successful approaches.

Operate security at the new speed of attack

Within two years, experts expect AI to favor attackers over defenders by 31.7%.[1] This projection highlights the growing threat posed by AI-driven attacks and the need for speed in security. That speed can be gained by using AI and agents. While most organizations apply AI only in limited areas of security operations—most commonly in threat detection—they should expand that use comprehensively across every aspect of security operations.

A good place to start is by applying agentic AI to vulnerability management. This area has become a soft target because of the capabilities of frontier AI models for identifying critical vulnerabilities in major operating systems and software browsers. By integrating AI across the security lifecycle, enterprises can transform their security programs into cohesive systems capable of continuous, autonomous defense.

This approach involves leveraging AI to analyze exposures, enforce policies, and coordinate detection and containment with minimal human intervention, reducing exposure windows and accelerating the containment of high-velocity attacks. By embracing these AI-driven security measures, enterprises can enhance their resilience, streamline compliance and ensure they're well-equipped to confront AI-powered threats at speed.

Shift identity security to continuous, runtime verification

Identity security remains underfunded despite mounting evidence of its criticality. Among organizations that suffered AI-related breaches in this year's study, 92% lacked proper access controls. This finding reveals a systemic failure to treat identity as mission-critical infrastructure. In the race to deploy AI agents to automate business processes—including security operations defending against frontier AI models—teams must fundamentally transform identity systems to secure not just humans, but NHIs, too.

This identity transformation will require:

- Agent discovery and onboarding
- Authentication delegation and authorization
- Anomaly detection and risk remediation
- Governance and lifecycle controls

This transformation will replace legacy approaches with modern, zero trust principles. Identity systems can then help ensure AI agents operate safely within intended guardrails—enforced through just-in-time access, time-bound approvals, and continuous risk-based runtime controls from devices, users and workloads. Identity security, which can be strengthened by autonomous security tools and services, must operate with no friction and at machine speed. Organizations making this shift can safely adopt AI at scale, meeting regulatory requirements and gaining significant competitive advantages.

Strengthen AI control through AI sovereignty

As organizations scale AI across platforms, cloud and ecosystems, they can lose control over how models operate and how data is transformed and accessed. This challenge increases the risk of breaches and responses to them. That's why AI sovereignty—maintaining control over where AI systems run, how data is processed and who has access—is now a critical security imperative.

AI sovereignty should include a security posture that extends across data, applications, identities and cloud environments where AI runs and is used. This security should be comprehensive, not piecemeal. It should include three dimensions:

strengthening app and API security, enforcing trusted identity controls for users, data, and machine agents, and ensuring cloud workloads are configured and monitored.

Monitoring how data enters, transforms within and exits AI systems can help organizations proactively identify sensitive data exposure risks, strengthen AI governance and compliance, and securely scale AI adoption.

By unifying these layers—supported by integrated approaches—organizations can better reduce systemic risk and build confidence in deploying AI at scale. Establishing control includes limiting unnecessary data movement, securing integrations with external tools and models, and maintaining clear visibility across workloads. It means strengthening app and API security and ensuring cloud workloads are configured and rigorously monitored.

Prepare for post-quantum security risks

Foundational gaps in encryption and cryptography management leave organizations exposed, even as investments in quantum-safe security have begun to rise. For instance, only 37% of breached organizations reported encrypting sensitive data at the time of breach. Simultaneously, just 34% reported having controls in place to monitor and secure cryptographic assets such as keys and certificates across their environments.

Cryptography touches every corner of the digital world, from internet protocols and enterprise applications to critical infrastructure and financial systems. To manage cryptography effectively in the quantum era, organizations should start transitioning to post-quantum encryption algorithms and modernizing their cybersecurity practices.

To accelerate the journey to post-quantum cryptography, establish crypto-agility and secure data across hybrid environments and AI, organizations must begin to build a robust security posture. This posture will allow organizations to discover cryptography in their environment, analyze vulnerabilities and remediate risks. The journey can be complex and dependent on each organization's needs. Therefore, it can be useful to have experts who understand the organization's unique needs and processes serve as a guide to quantum-safe migration.

Organization demographics

This year's study examined 602 organizations of various sizes across 16 countries and geographic regions and 17 industries.

This section explores the breakdown of organizations in the study by geography and industry and defines the industry classifications.

Geographic demographics

The 2026 study was conducted across 16 countries and geographic regions. For the second year the study included Benelux, the economic union of Belgium, the Netherlands and Luxembourg.

ASEAN is a cluster sample of organizations located in Singapore, Indonesia, Philippines, Malaysia, Thailand and Vietnam. Latin America is a cluster sample of organizations located in Mexico, Argentina, Chile and Colombia. Middle East is a cluster sample of organizations located in Saudi Arabia and the United Arab Emirates.

Distribution by sample or region

ASEAN	4%	Australia	5%
US	10%	Benelux	6%
India	9%	Canada	5%
Brazil	7%	LATAM	5%
UK	8%	South Korea	5%
Germany	8%	ASEAN	4%
Japan	7%	Italy	4%
Middle East	7%	South Africa	4%
France	6%		

Industry demographics

The selection of 17 industries has been consistent across multiple years of the study. This year, the top 4 industries—financial, industrial, professional services and technology—accounted for 46% of the 602 organizations studied.

Industry			
Financial	13%	Consumer	4%
Industrial	12%	Hospitality	4%
Services	10%	Media	4%
Technology	11%	Pharma	4%
Energy	8%	Education	2%
Public	7%	Research	2%
Communications	6%	Healthcare	2%
Transportation	6%	Entertainment	1%
Retail	4%		

Industry definitions

Healthcare
Hospitals and clinics

Financial
Banking, insurance and investment companies

Energy
Oil and gas companies, utilities and alternative energy producers and suppliers

Pharmaceuticals
Pharmaceutical companies, including biomedical life sciences

Industrial
Chemical processing and engineering, and manufacturing companies

Technology
Software and hardware companies

Education
Public and private universities and colleges, and training and development companies

Professional services
Services such as legal, accounting and consulting firms

Entertainment
Movie production, sports, gaming and casinos

Transportation
Airlines, railroads and trucking, and delivery companies

Communications
Newspapers, book publishers, and public relations and advertising agencies

Consumer
Manufacturers and distributors of consumer products

Media
Television, satellite, social media and internet

Hospitality
Hotels, restaurant chains and cruise lines

Retail
Brick and mortar and e-commerce

Research
Market research, think tanks, and research and development

Public
Federal, state and local government agencies, and nongovernmental organizations

Research methodology

The numerical value obtained from the number line, rather than a point estimate for each presented cost category, preserved confidentiality and ensured a higher response rate.

The benchmark instrument also required respondents to provide a second separate estimate for indirect and opportunity costs.

In the interest of maintaining a manageable dataset for benchmarking, the report included only those cost activity centers with a crucial impact on data breach costs. Based on discussions with experts, a fixed set of cost activities was chosen. After collecting benchmark information, each instrument was carefully reexamined for consistency and completeness.

The scope of data breach cost factors was limited to known categories that apply to a broad set of business operations involving personal information. We chose to focus on business processes instead of data protection or privacy compliance activities because we believed the process study would yield better-quality results.

How we calculate the cost of a data breach

To calculate the average cost of a data breach, we excluded very small and very large breaches. Data breaches examined in the 2026 report ranged in size between 2,590 to 115,380 compromised records.

We used activity-based costing, which identifies activities and assigns a cost according to actual use. Four process-related activities drove a range of expenditures associated with an organization's data breach: detection and escalation, notification, post-breach response and lost business.

Detection and escalation
Activities that enable an organization to detect the breach include:

- Forensic and investigative activities
- Assessment and audit services
- Crisis management
- Communications to executives and boards

Notification
Activities that enable an organization to notify data subjects, data protection regulators and other third parties include:

- Emails, letters, outbound calls or general notices to data subjects
- Determination of regulatory requirements
- Communication with regulators
- Engagement of outside experts

Post-breach response
Activities to help victims of a breach communicate with an organization and conduct redress activities to victims and regulators include:

- Help desk and inbound communications
- Credit monitoring and identity protection services
- Issuing of new accounts or credit cards
- Legal expenditures
- Product discounts
- Regulatory fines

Lost business
Activities that attempt to minimize the loss of customers, business disruption and revenue losses include:

- Business disruption and revenue losses due to system downtime
- Cost of losing customers and acquiring new customers
- Reputational damage and diminished goodwill

Data breach FAQs

What's a data breach?
A data breach is defined as an event in which records containing PII; financial or medical account details; or other secret, confidential or proprietary data are potentially put at risk. These records can be in electronic or paper format. Breaches included in the study ranged between 2,590 to 115,380 compromised records.

What's a compromised record?
A record is information that reveals confidential or proprietary corporate, governmental or financial data, or identifies an individual whose information has been lost or stolen in a data breach. Examples include a database with an individual's name, credit card information and other PII, or a health record with the policyholder's name and payment information.

How do you collect the data?
Our researchers collected in-depth qualitative data over 3,558 separate interviews with individuals at 602 organizations that suffered a data breach between March 2025 through February 2026. Interviewees were familiar with their organization's data breach and the costs associated with resolving the breach. These interviewees included CEOs or executives, heads of operations, controllers or heads of finance, IT practitioners, business unit leaders and general managers, and risk management and cybersecurity practitioners. For privacy purposes, we didn't collect organization-specific information.

For the first time, researchers went back to these same organizations, in May 2026, following news that frontier AI models had developed advance threat capabilities and asked if these new threats had changed their security spending plans. Among the original 602 organizations, 456 responded and 78% of those that did said they were aware of these new threats.

What's included in the cost of a data breach?
We collected both the direct and indirect expenses incurred by the organization. Direct expenses included engaging forensic experts, outsourcing hotline support and providing free credit monitoring subscriptions and discounts for future products and services. Indirect costs included in-house investigations and communications along with the extrapolated value of customer loss resulting from turnover or diminished customer acquisition rates.

This research represented only events directly relevant to the data breach experience. Regulations, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), may encourage organizations to increase investments in their cybersecurity governance technologies. However, such activities didn't directly affect the cost of a data breach for this research. For consistency with prior years, we used the same currency translation method rather than adjusting accounting costs.

How does benchmark research differ from survey research?
The unit of analysis in the Cost of a Data Breach Report was the organization. In survey research, the unit of analysis is the individual. We recruited 602 organizations to participate in this study.

Can the average per-record cost be used to calculate the cost of breaches involving millions of lost or stolen records?
It's not consistent with this research to use the overall cost per record as a basis for calculating the cost of single or multiple breaches totaling millions of records. The per-record cost is derived from our study of hundreds of data breach events in which each event featured a maximum of 115,380 compromised records.

Are you tracking the same organizations each year?
Each annual study involves a different sample of organizations. To be consistent with previous reports, we recruit and match organizations each year with similar characteristics, such as the organization's industry, head count, geographic footprint and size of data breach. Since starting this research in 2005, we have studied the data breach experiences of 7,087 organizations.

Research limitations
Our study used a confidential and proprietary benchmark method that was successfully deployed in earlier research. However, the inherent limitations with this benchmark research need to be carefully considered before drawing conclusions from findings.

Nonstatistical results
Our study drew upon a representative, nonstatistical sample of global entities. Statistical inferences, margins of error and confidence intervals can't be applied to this data, given that our sampling methods weren't scientific.

Nonresponse
Nonresponse bias wasn't tested, so it's possible that organizations that didn't participate are substantially different in terms of underlying data breach cost.

Sampling-frame bias
Because our sampling frame was judgmental, the quality of results was influenced by the degree to which the frame was representative of the population of organizations being studied. We believe the current sampling frame was biased toward organizations with more mature privacy or information security programs.

Organization-specific information
The benchmark didn't capture organization-identifying information. Individuals could use categorical response variables to disclose demographic information about the organization and industry category.

Unmeasured factors
We omitted variables from our analyses, such as leading trends and organizational characteristics. The extent to which omitted variables might explain benchmark results can't be determined.

Extrapolated cost results
Although certain checks and balances can be incorporated into the benchmark process, it's always possible respondents didn't provide accurate or truthful responses. In addition, the use of cost extrapolation methods rather than actual cost data may inadvertently introduce bias and inaccuracies.

Currency conversions
The conversion from local currencies to the US dollar deflated average total cost estimates in other countries. For purposes of consistency with prior years, we decided to continue to use the same accounting method rather than adjust the cost. It's important to note this issue may affect only the global analysis because all country-level results are shown in local currencies.

The current real exchange rates used in this research report were published by the Federal Reserve on 1 March 2026.

About us

Ponemon Institute

Founded in 2002, Ponemon Institute is dedicated to independent research and education that advances responsible information and privacy management practices within business and government. Ponemon's mission is to conduct high-quality empirical studies on critical issues affecting the management and security of sensitive information about people and organizations.

Ponemon Institute upholds strict data confidentiality, privacy and ethical research standards and doesn't collect any personally identifiable information (PII) from individuals or company-identifiable information in business research. Furthermore, strict quality standards ensure subjects aren't asked extraneous, irrelevant or improper questions. If you have questions or comments about this research report, including requests for permission to cite or reproduce the report, contact us by letter, phone call or email:

Ponemon Institute LLC
Research Department
1-800-887-3118
research@ponemon.org

IBM

IBM is a leading hybrid cloud, AI and business services provider, helping clients in more than 175 countries capitalize on insights from their data, streamline business processes, reduce costs and gain the competitive edge in their industries. All of it is backed by IBM's legendary commitment to trust, transparency, responsibility, inclusivity and service. For more information, visit ibm.com.

Learn more about advancing your security posture: visit ibm.com/security.

Join the conversation in the IBM Security Community.

1. Frontier AI's Impact on the Cybersecurity Landscape, UC Berkeley, 27 November 2025.