cyberhaven

2026 AI Adoption & Risk Report

Cyberhaven Labs

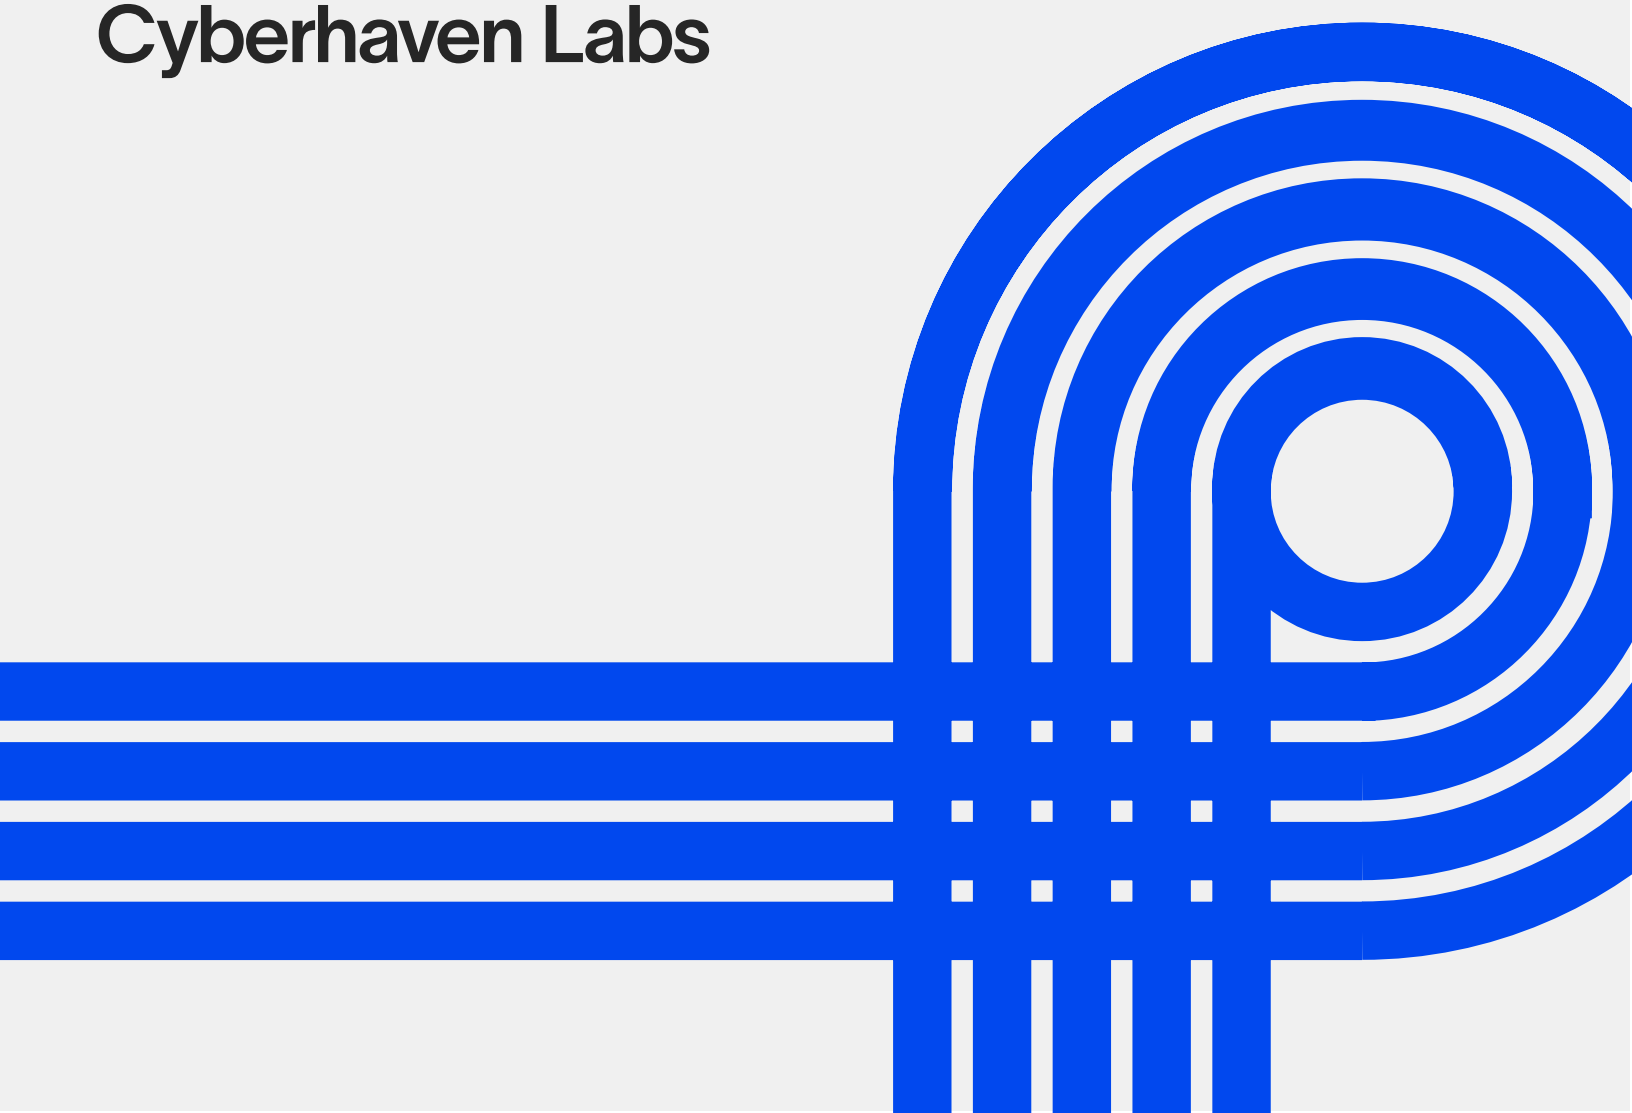

Table of Contents

Introduction

Since the launch of ChatGPT in 2022, AI has become one of the fastest-adopted workplace technologies in history.

What began as individual employees experimenting with generative AI has rapidly evolved into tools embedded directly into core business workflows across organizations of all sizes. The speed of this transition, from novelty to operational dependency, has been unprecedented. In many cases, it has outpaced the ability of enterprises to understand, govern, and secure AI usage.

AI adoption has not progressed evenly across the enterprise. Usage is increasingly polarized. A small but growing set of organizations is moving aggressively, deploying dozens or even hundreds of AI tools across development, operations, and knowledge work. Others remain cautious or only lightly engaged. In high-adoption environments, innovation often advances faster than governance. Experimentation frequently takes priority over visibility, control, and risk management.

This year's research from Cyberhaven Labs captures a clear shift in how enterprises are using AI and where risk is concentrating as a result.

To understand the full scope of adoption, we analyzed AI usage across three categories: Generative AI SaaS applications, endpoint AI applications, and AI agents. Drawing on billions of real-world data movements from hundreds of thousands of employees at a sample of 222 companies, we measured adoption using active user counts and event-level activity. This approach allowed us to assess not just whether AI is present, but how deeply it is embedded into daily work.

The data shows that while usage of traditional chat-based GenAI SaaS tools is beginning to plateau, enterprise AI adoption is not slowing down. In 2025, AI coding assistants, browser-based agents, and custom AI agents saw rapid growth. This second wave of adoption is more operational and more automated. It is also far more difficult to govern. These tools operate inside development environments, browsers, and workflows. They interact directly with sensitive data, proprietary code, and critical systems, often with limited oversight.

As AI becomes infrastructure rather than a standalone interface, the security implications intensify. Employees are no longer using AI only for ideation or research. They are inputting source code, financial data, customer information, and intellectual property across a fragmented and expanding ecosystem of tools. Much of this activity occurs outside traditional IT visibility. It spans personal accounts, open-weight models, and SaaS platforms that lack enterprise-grade security controls.

The result is a familiar pattern, amplified by scale and speed. Shadow adoption increases. Controls are applied inconsistently. Risk accumulates faster than most organizations can measure or manage it. The risks associated with enterprise AI use are no longer theoretical or future-dated. They are already material, unevenly distributed, and concentrated among the organizations and teams adopting AI most aggressively.

This report provides a data-driven view into how enterprises used AI in 2025, where adoption is accelerating, and where security risk is compounding. By examining real-world usage patterns across industries, departments, tools, and data types, Cyberhaven Labs aims to help security and technology leaders understand not just the scale of AI adoption, but the context required to govern it safely as they plan for 2026.

Key Findings

01 Organizations with the highest rates of AI adoption are utilizing **over 300** GenAI tools within their enterprise environment.

02 Chinese open-weight models are now enterprise favorites, accounting for **50%** of endpoint-based usage among Cyberhaven users.

03 GenAI tools remain risky across the board. When looking at the top 100 most-used GenAI SaaS applications, **82%** are classified as "medium," "high," or "critical" risk.

04 **One-third** of employees are accessing GenAI tools from personal accounts, increasing overall risk and Shadow AI.

05 Employees are feeding AI tools sensitive data, as over a third (**39.7%**) of all interactions with AI tools involve sensitive data.

1

An AI Adoption Gap is Emerging

Artificial intelligence (AI) and large language models (LLMs) are becoming increasingly embedded in organizational workflows.

Today, 62% of organizations[1] are experimenting with AI agents, enterprises are spending four times more on AI software than on traditional software, and 74% of executives stated[2] they achieve returns within the first year of AI tool deployment.

However, AI adoption and use is not unfolding as a steady, industry-wide wave. Instead, it is becoming increasingly polarized.

A widening gap is emerging between AI early adopters and organizations that remain hesitant to embrace these technologies.

Frontier enterprises — those with the highest rates of AI adoption — are interacting with hundreds of GenAI applications over the course of 2025. In the most advanced cases, organizations are using more than 300 GenAI tools, while even the broader group of AI leaders averages over 200. By contrast, cautious enterprises typically employ fewer than 15 GenAI tools.

This divide is stark when compared to the median organization, which uses 54 GenAI applications. In practice, frontier enterprises are adopting AI tools at nearly six times the rate of the average company.

1 https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

2 https://blog.arcade.dev/ai-integration-platform-trends

GenerativeAI SaaS

The same divide appears at the employee level.

In the average organization, roughly one-third of employees use GenAI tools regularly. Yet adoption rates vary dramatically by enterprise maturity. Frontier organizations see a 71.4% employee adoption rate, while the most cautious enterprises report adoption as low as 2.5%. As organizations deploy more GenAI tools, employee usage predictably mirrors that investment.

As discussed in Section 3, industry-level variance likely contributes to this wide distribution, with certain sectors accounting for a disproportionate share of GenAI usage while others lag behind.

Distribution of Number of GenAI SaaS Applications used by an Enterprise

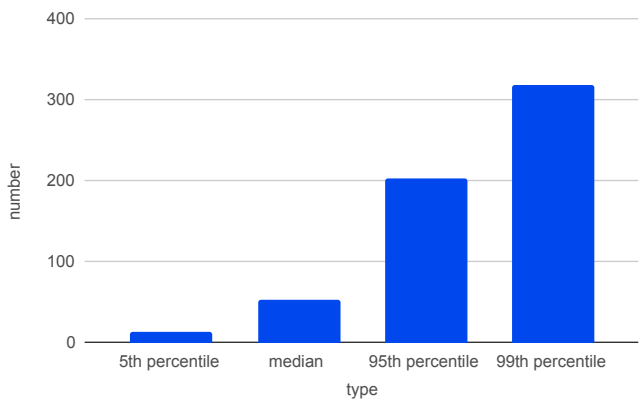

Percentage of Employees using GenAI SaaS

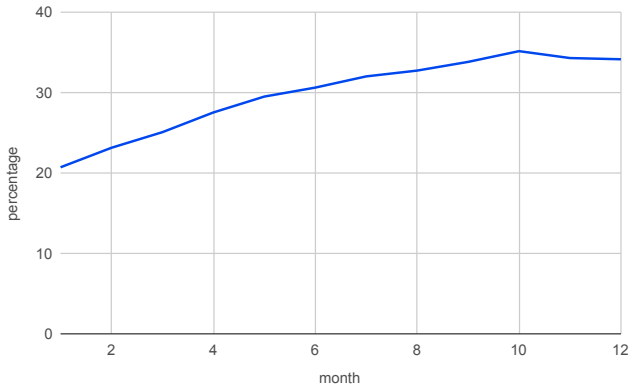

Most Organizations Remain Hesitant to Adopt AI

While frontier organizations are rapidly experimenting with GenAI, the majority of enterprises remain cautious. Within the median organization, only 33.4% of employees have adopted AI tools, and this broad hesitation is a key driver of today's uneven adoption landscape.

This polarized adoption pattern reveals two realities. Some organizations are aggressively adopting AI and may realize outsized gains in innovation and growth. At the same time, these frontier enterprises are also assuming a disproportionate share of AI risk.

As data flows through hundreds of GenAI tools, rapid adoption multiplies risk points, governance complexity, and potential sensitive data exposure. Many organizations appear to be trading coordination and security controls for experimentation, creating a growing gap between AI adoption and AI security. This challenge is further amplified by uneven employee adoption rates, making one-size-fits-all AI security approaches ineffective. Effective AI security will depend not only on which tools are deployed, but on how — and by whom — they are actually used.

2

Enterprise AI Usage Spans Models, Tools, and Accounts — Outpacing Governance

Chinese Open-Weight Models Growing in Popularity

While AI is often synonymous with "ChatGPT" in popular discourse, much as "Google" became shorthand for search, the Silicon Valley models are not the only ones seeing widespread use. Chinese open-weight models such as Qwen and DeepSeek have maintained a steady share of workplace AI adoption after the initial wave of hype, raising ongoing governance and national security considerations for global enterprises.

When looking at Cyberhaven's userbase, endpoint-based Chinese open-weight models account for 50% of usage. Usage spiked early in the year, driven largely by DeepSeek's rapid rise, before stabilizing. Despite the waning of initial momentum, Chinese models account for 50% of endpoint-based open-weight model usage.

When we widen this lens, open-weight models in general are starting to make waves in the market. According to a recent report by OpenRoute[3], open-weight model usage accounted for about 30% of all usage, and over one year Chinese open-weight models averaged 13% of all AI usage, including both open-weight models and propriety models.

Percentage of Chinese Open Weight Model Usage Over Time

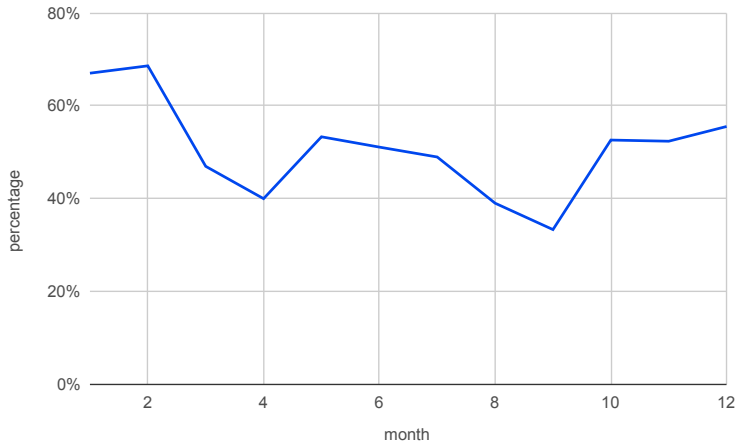

3 https://openrouter.ai/state-of-ai

For enterprises, these models can introduce elevated risk. Key risk factors include:

Regulatory and legal exposure, as Chinese AI developers can operate under national security, intelligence, and data laws that may compel cooperation with state authorities

Supply chain and governance opacity, with limited visibility into training data sources, data provenance, fine-tuning practices, and post-release controls

Amplified data sovereignty and leakage risks, particularly when models are hosted, updated, or integrated through infrastructure tied to Chinese vendors or ecosystems

At the same time, enterprises and individuals continue to adopt a wide range of GenAI tools, underscoring how quickly the AI market has become saturated and fragmented.

Closed-Weight vs. Open-Weight AI

Open and closed AI models present different tradeoffs for enterprises.

While **open-weight models** can offer flexibility and cost advantages, they often introduce greater governance, provenance, and security challenges.

Closed models may provide stronger controls and contractual assurances, but limit transparency and customization.

Effective AI governance requires understanding and managing these tradeoffs.

Top 20 GenAI SaaS Applications in 2025

1. chatgpt.com
2. gemini.google.com
3. claude.ai
4. grok.com
5. perplexity.ai
6. copilot.microsoft.com
7. lovable.dev
8. chat.deepseek.com
9. app.glean.com
10. notebooklm.google.com
11. app.devin.ai
12. app.harvey.ai
13. app.blueflame.ai
14. doubao.com
15. midjourney.com
16. aistudio.google.com
17. poe.com
18. otter.ai
19. app.codesignal.com
20. ai.azure.com

Tool Usage and Prominence Continues To Fluxuate

Compared to 2024, the top GenAI SaaS list includes new entrants such as Grok, Lovable, and Blueflame. Notably, Gemini has surpassed Claude in total event volume — after Claude held second place in 2024 — and Google now accounts for three of the top 20 GenAI applications overall, signaling its rapid emergence as a major AI platform provider.

At the same time, several tools that appeared in prior years' rankings, including Phind, Jasper.ai, Copy.ai, Dialogflow, and TensorFlow, have dropped off the list entirely. This churn highlights the continued volatility of the GenAI market. While established platforms like Gemini, Claude, and Copilot may approach the ubiquity of tools like PowerPoint, smaller vendors continue to rise and fall quickly.

This diversity also underscores a critical security reality: GenAI tools vary widely in safeguards, maturity, and governance capabilities. Organizations must evaluate each tool as distinct software with its own risk profile, rather than treating AI as a single, uniform category.

Most GenAI SaaS Tools Are Objectively Risky

When GenAI tools are evaluated by risk level, the results are stark.

Across the top 100 most-used GenAI SaaS applications, 82% are classified as "medium," "high," or "critical" risk. Even when excluding "medium" risk and considering only "high" and "critical," 60% of tools still fall into these categories.

For security leaders, this means that most AI usage today occurs in tools that would not meet traditional enterprise risk standards — yet employees continue to input sensitive data into them at high rates (see Section 4).

RiskIQ distribution among Top 100 most used GenAI SaaS

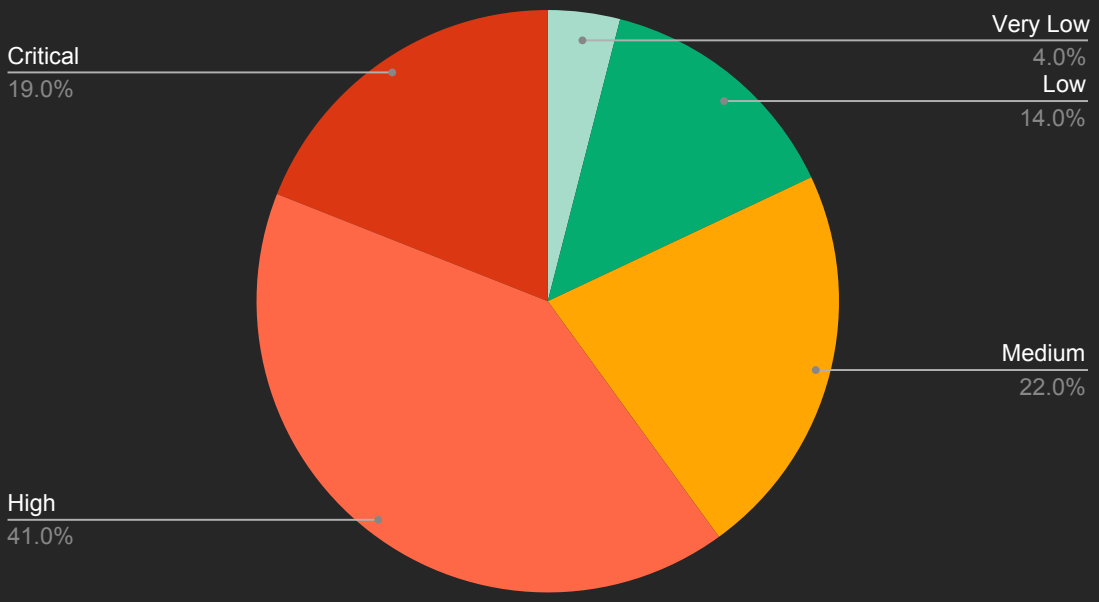

The risk is determined by the Cyberhaven AI App RiskIQ engine, which is a deep research based AI agent.

We evaluate AI App risk on 5 categories:

- Data sensitivity and Security
- Model Security Risks
- Compliance & Regulatory Risks
- User Authentication & Access Controls
- Security Infrastructure & Practices

To come up with a comprehensive risk level

Personal Account Usages Creates Risk

Not all employees access GenAI tools through governed corporate accounts. Cyberhaven's browser extension uniquely identifies which cloud account an employee uses when interacting with AI applications.

Our data shows that 32.3% of ChatGPT usage occurs through personal accounts, as does 24.9% of Gemini usage. Claude and Perplexity see even higher rates of personal account usage, at 58.2% and 60.9% respectively. This behavior significantly limits organizational visibility into AI usage and data flows.

Percentage of Personal Account Use Per Platform

32.3%	24.9%	58.2%	60.9%
ChatGPT	Gemini	Claude	Perplexity

Personal account usage reduces control and oversight, increasing the likelihood of data policy violations, data leakage, and compliance risk. Employees may use personal accounts for convenience, through negligence, or to bypass corporate restrictions — but these workarounds create measurable security exposure. The risk intensifies when weak governance exists at multiple layers. Tools that lack enterprise-grade controls — including some open-weight and non-U.S. models — are often accessed through personal accounts, creating overlapping blind spots in both model governance and user access.

Addressing this risk requires both employee education and stronger mechanisms to detect, govern, and manage shadow AI usage.

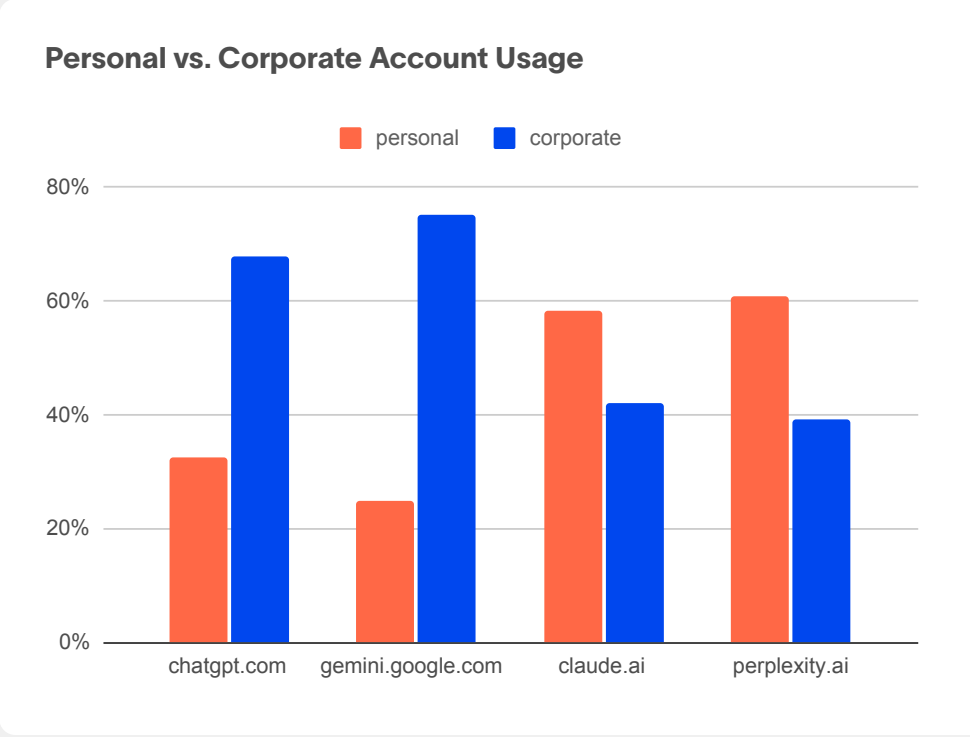

3

Regulated Industries Are Adopting AI Most Frequently

AI adoption varies significantly by industry. The technology sector leads, with 40.5% of employees using AI tools, followed by pharmaceuticals at 33% and financial services at 28.7%.

This distribution is unsurprising. Technology organizations have historically been early adopters of productivity and operational tools and tend to favor experimentation, even when it means accepting a higher risk tolerance.

The pharmaceutical industry is increasingly turning to AI to drive innovation. Facing persistent budget and labor constraints, along with rising consumer expectations fueled by telehealth and digital-native pharmaceutical companies, the industry is under pressure to accelerate the digitization of both operations and customer experience (e.g chatbots, more personalized engagement, simple digital portals, and predicting customer needs).

Financial services show a similar pattern. Institutions are adopting AI to streamline core business processes—such as delivering highly customized portfolio recommendations to clients—while also enhancing customer-facing experiences and service efficiency.

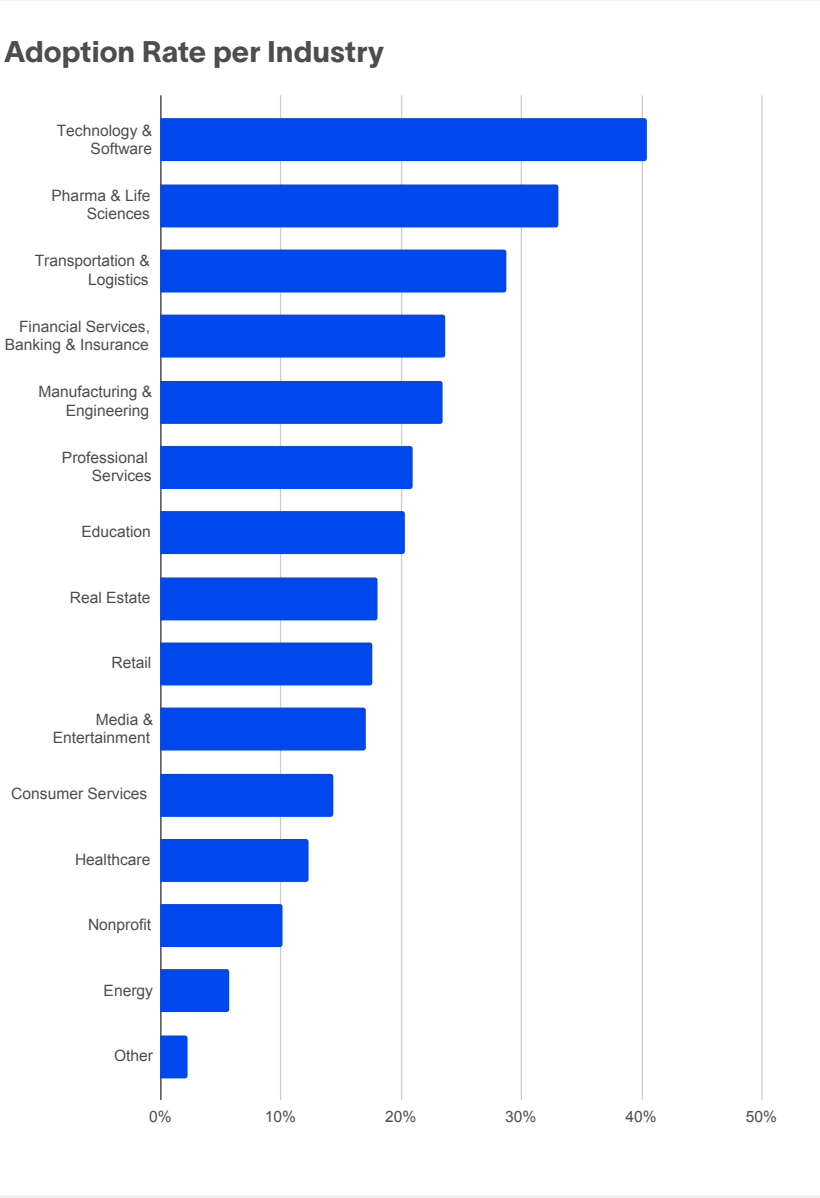

Department By Department Usage Skews Toward Technical Teams

Within organizations, **AI usage is highest in engineering departments**, where more than 60% of employees use AI tools—nearly 20 percentage points higher than the next highest group, marketing.

This gap reflects engineers' propensity to adopt new technologies and the growing role of AI in supporting engineering workflows, from automating routine coding tasks to assisting with complex IT and problem-solving challenges.

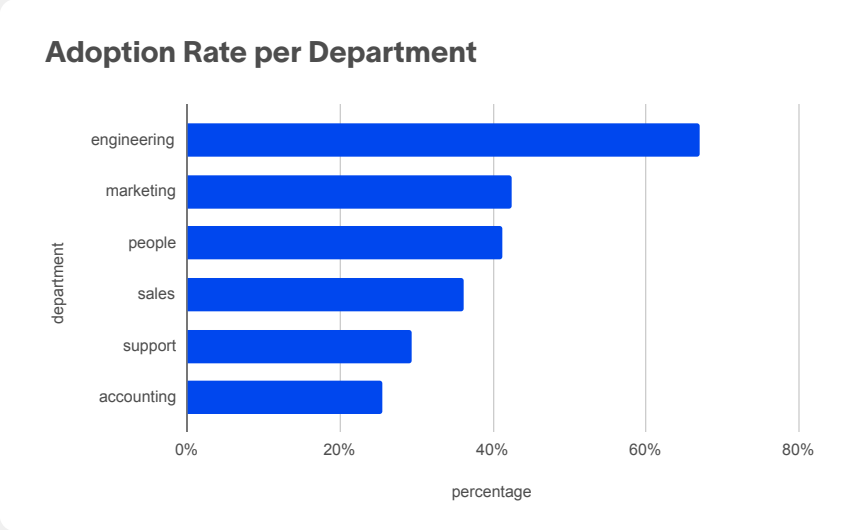

MacOS Preferred by Engineers

When comparing the aggregate of all departments, we see that half – 55% – of departments use MacOS for these AI agents. **For engineers, that percentage jumps to 63%.**

These adoption patterns carry important governance implications. Industries with the highest levels of AI usage, such as pharmaceuticals and financial services, also tend to manage highly regulated and sensitive data, raising the stakes for AI oversight and control. At the same time, AI usage concentrated in engineering teams signals that these tools are being embedded directly into core systems and workflows, rather than remaining a surface-level productivity enhancement.

4

Employees Are Pouring Sensitive Data into Risky AI Tools

AI usage is growing rapidly across organizations, but this growth comes with new security risks.

Employees are not just asking questions of GenAI tools; they are inputting data – often fragmented and moved across systems – directly into these technologies.

From source code to drug research, financial information, and intellectual property, GenAI is handling some of the most sensitive data in organizations today.

Categorization of Data put into AI (over all industries)

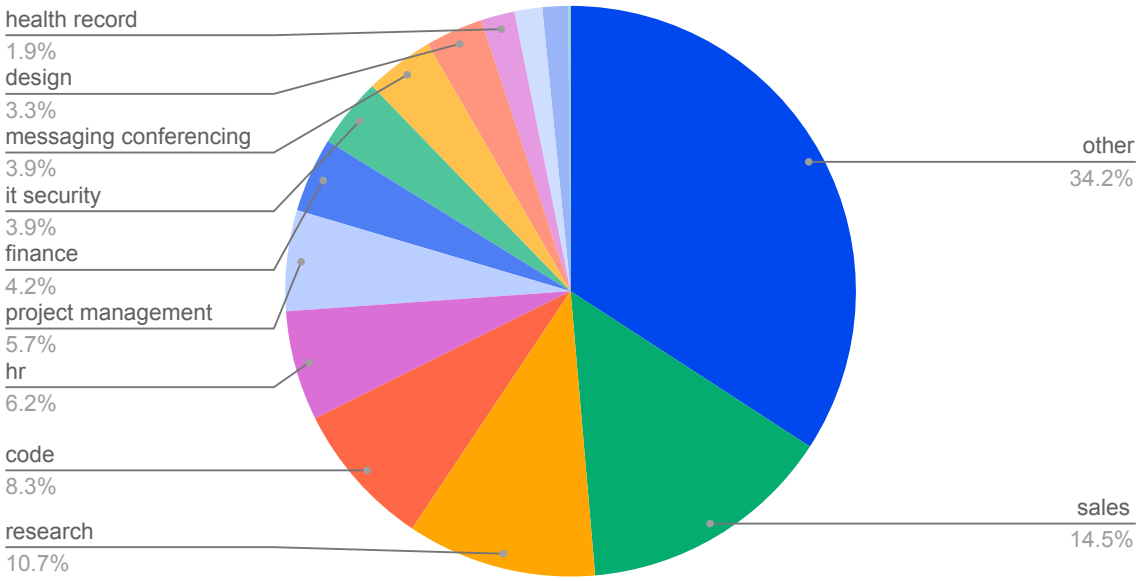

Sensitive Data, Defined

Sensitive data can be subjective and depends on an organization's labeling rules. For the purposes of this report, Cyberhaven defines sensitive data as any data marked "sensitive" per customer-set rules, except for customers who classify all data as sensitive. Common examples include personally identifiable information (PII), personal health data (PHI), financial data, customer data, and business-specific data.

A concerning amount – 39.7% of all interactions with AI tools – involve sensitive data, including prompts or copy-paste actions. That means the average employee enters sensitive data into AI tools once every three days.

The data comes from multiple parts of the business, not just isolated workflows:

- **Sales and go-to-market data:** mid-teens percentage of AI-bound data globally and nearly 30% of what sales teams send into AI
- **Research and R&D content:** the dominant category in industries like healthcare, where about one-third of AI-bound data is research-related
- **Technical and other controlled assets:** source code, internal project data, and sensitive categories like health records that traditionally require tight controls

This trend is not driven by malicious behavior. Employees are typically seeking productivity gains, faster answers, or ways to solve complex problems—all benefits that AI tools provide.

However, without proper governance, monitoring, and controls around both the tools and the data inputted into them, organizations face serious risks: data leakage, shadow AI growth, and exposure from tools storing or reusing sensitive information across users.

5

Coding Assistants and AI Agents are Becoming The "Second Wave" of Workplace AI

Workplace AI is entering a “second wave,” moving beyond general-purpose tools to more specialized applications that directly enhance workflows. Coding assistants and AI agents are no longer niche experiments — they are rapidly becoming embedded in the daily operations of developers and teams across enterprises. This shift reflects both the maturation of AI adoption and the growing appetite for tools that increase productivity, automate complex tasks, and integrate directly into core systems.

AI Coding Assistants See Same Adoption Gap

AI coding assistants (such as Cursor, GitHub Copilot, and Claude Code) continued steady growth through 2025.

Key trends:

- Roughly **half of all developers (49.5%)** were using coding assistants by December, up from about 20% at the start of the year.
- In leading companies, nearly **90% of developers** use these tools, while in a typical organization, adoption is closer to 50%. This illustrates the **growing gap of AI adoption**, with developers at frontier companies being **11.5× more likely** to use AI coding assistants.
- Developers are increasingly using multiple assistants: the share using **two or more** doubled from 16% in January to 32% in November.

These trends suggest that AI is not just a productivity enhancer but is increasingly embedded into core development workflows, mirrored by the data that technology departments are using AI most frequently. This is a shift that raises security considerations. As developers input source code and internal project data into these tools, organizations must evaluate governance, monitoring, and access controls to prevent sensitive data leakage.

The same focus on specialized workflows is driving enterprises to adopt custom AI agents and platforms, where early adoption is concentrated in highly regulated and operationally complex industries.

Percentage of Developers Using AI Coding Assistants

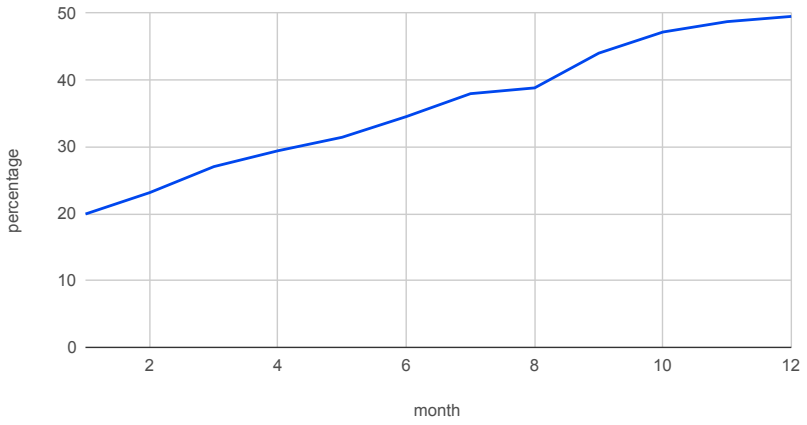

Distribution of Developers using AI across companies

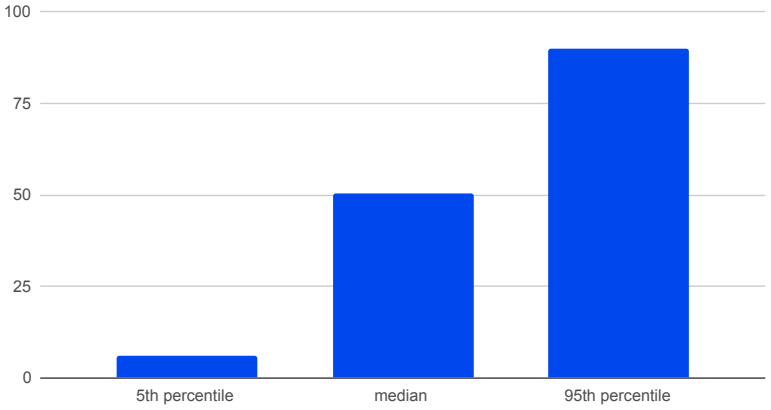

Percentage of users per AI coding assistant

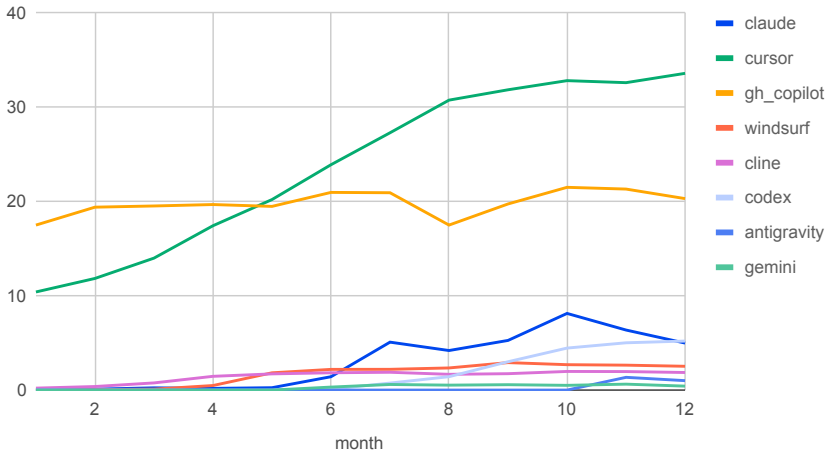

Organizations Are Investing In Custom Agents and Workflows

Nearly 23% of enterprises have adopted agent-building platforms (such as n8n, OpenAI Agent Builder, Glean Agents, or Microsoft Copilot Studio) to create custom AI agents and workflows.

Adoption is concentrated in pharmaceutical, manufacturing, and energy companies, suggesting early uptake is clustering in industries that are both highly regulated and operationally complex. This mirrors trends in GenAI adoption: while technology leads in GenAI tool adoption, pharmaceutical companies are now leading in AI agent use.

Pharmaceutical companies leverage these custom-built AI agents to address workforce and skill shortages, bridge operational gaps, and enhance systems, all while navigating strict compliance requirements and regulatory oversight. In this way, the industry is harnessing AI to drive innovation without sacrificing security.

Importantly, as organizations embed AI agents into core business processes, sensitive data from operations, research, and workflows is increasingly interacting with these tools, amplifying the need for monitoring, governance, and secure access controls. In this way, the industry is harnessing AI to drive innovation without sacrificing security.

Percentage of Companies using Agent building SaaS Platforms

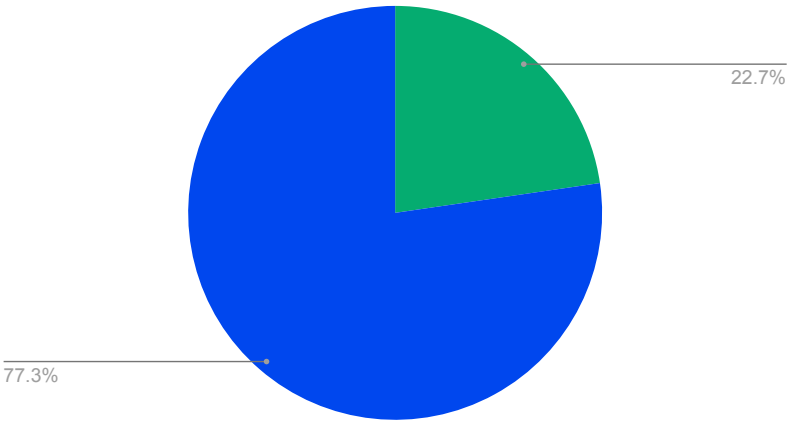

Percentage of companies by Industry

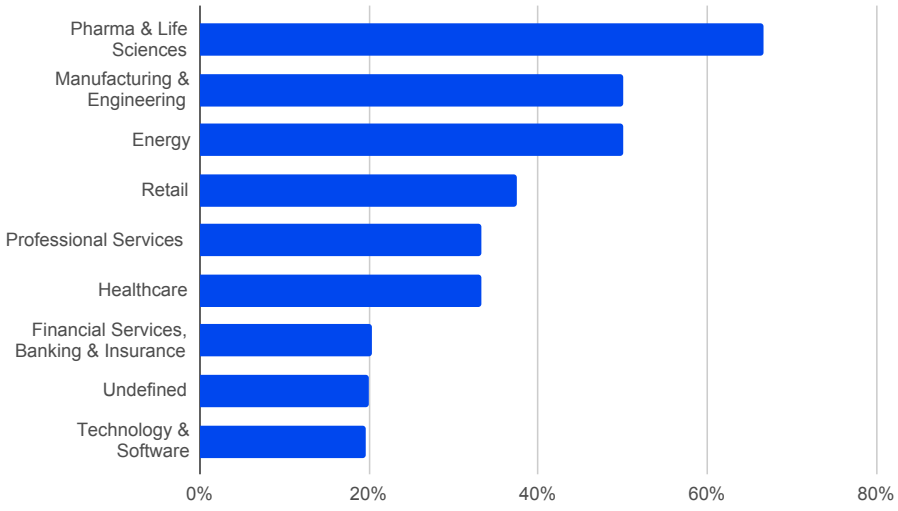

Conclusion:

AI Security Is Paramount As Organizations Race To Adopt New Technology

AI adoption is accelerating, but it is not unfolding evenly. This research shows that enterprise AI usage is deeply polarized, with a small group of frontier organizations driving outsized adoption while others remain cautious. In many cases, growth and experimentation are prioritized first, while security, governance, and oversight follow later. At the same time, most AI usage today occurs in tools that carry elevated risk, and employees are routinely inputting sensitive data into a wide and growing ecosystem of GenAI tools, coding assistants, and custom-built agents.

For many organizations, this moment represents a kind of "wild west" for AI. Tools are proliferating faster than policies, employee usage often outpaces visibility, and sensitive data flows across models, applications, and accounts with limited centralized control. Without closer scrutiny of how AI is actually being used — by which teams, in which workflows, and with what data — organizations risk creating an expanding gap between innovation and security.

As AI adoption becomes more uneven across industries, departments, and individual users, AI security must become a top priority. One-size-fits-all policies are unlikely to succeed in environments where some teams rely heavily on AI while others barely use it at all. Effective governance requires understanding real usage patterns and applying controls that reflect the sensitivity of the data, the maturity of the user group, and the risk profile of the tools involved.

AI security is also inherently complex. The scale, speed, and fragmentation of AI adoption make it difficult for organizations to manage risk through manual processes or isolated controls alone. For many enterprises, partnering with specialized third-party security providers can help simplify this challenge — bringing visibility, context, and enforcement together across data, users, and AI systems. As organizations continue to race toward AI-driven innovation, those that invest early in comprehensive AI security will be best positioned to move fast without compromising trust, compliance, or resilience.

In this environment, approaches grounded in data security posture management (DSPM) and data lineage are increasingly critical. By providing continuous visibility into where sensitive data lives, how it moves, and which AI tools and users interact with it, organizations can move from reactive enforcement to proactive, context-aware governance that scales with AI adoption.

Cyberhaven And the Future of AI Security

Enterprises face a growing blind spot: shadow AI. Employees adopt generative AI tools on their own, often sharing sensitive data without oversight. Cyberhaven's AI Data Security solution changes that. We discover AI tools in use, provide detailed risk assessments, and enforce controls on data flows to and from AI. This lets companies stop leaks to risky tools, protect confidential data, and enable safe adoption of AI at scale. Experience unprecedented visibility and protection for AI usage with Cyberhaven.