

2026 AI Adoption and Risk Report

Six-Month Update

New data on how enterprise AI adoption shifted from December 2025 through June 2026, and what it means for data security in the agentic enterprise.

Introduction

In January 2026, Cyberhaven Labs published the 2026 AI Adoption & Risk Report, documenting how generative AI (GenAI) software as a service (SaaS), endpoint AI applications, and AI agents were reshaping enterprise workflows. Six months of new telemetry now let us test those findings against what shifts occurred in real-time.

The picture that emerges is a story of AI adoption moving from applications to workflows. Chat-based GenAI SaaS tools are consolidating and plateauing, even as the volume of data moving through them grows sharply. Meanwhile, endpoint agentic AI, including coding assistants, autonomous agents, and desktop AI applications, is accelerating past GenAI SaaS in both adoption rate and security relevance.

METHODOLOGY

This update applies the same Cyberhaven Labs data methodology as the original 2026 AI Adoption & Risk Report: real-world data movement events and active user counts across Cyberhaven’s enterprise customer base, covering December 2025 through June 2026.

BY THE NUMBERS

Key Takeaways

37.7% of employees now use GenAI SaaS applications, up from a 35% plateau that held from January through May 2026.	+80% year-over-year increase in data movement events into and out of GenAI SaaS, June 2025 to June 2026.	6.45 vs. 2.15 average daily Claude Code launches per machine vs. Codex, among companies with 100+ Claude Code users.
47.6 vs. 50.5 distinct GenAI SaaS applications used per enterprise in H1 2026 vs. H1 2025, a possible sign of tool consolidation.	60.5% peak developer adoption of AI coding assistants in May 2026, up from 49.5% in December 2025, before easing to 55.5% in June.	+43.2% growth in non-coding endpoint agentic app users (Claude Desktop, ChatGPT Desktop, MS Copilot) from January to June 2026, led by Claude Desktop’s 1,233% surge.

Codex +16 / Cursor -11

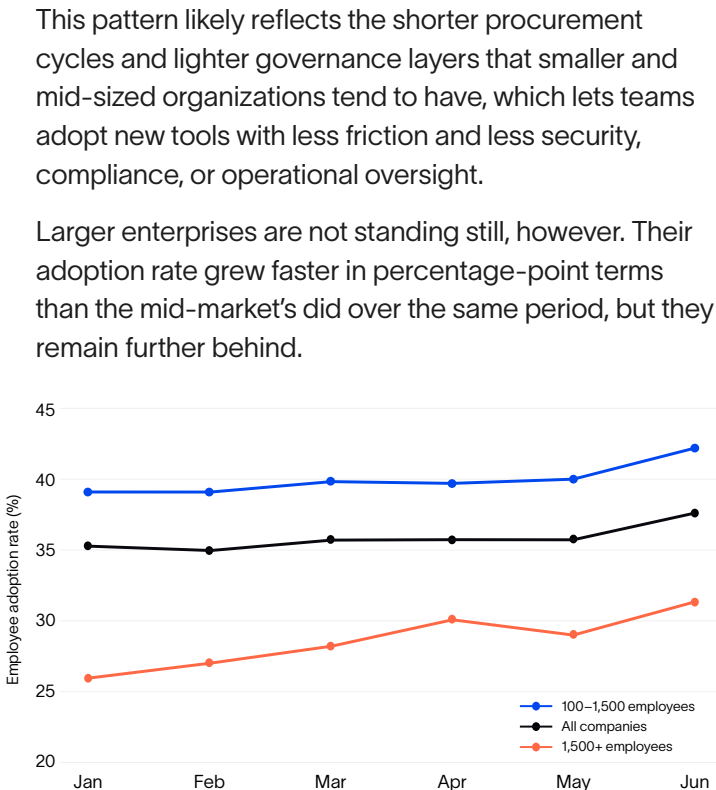
percentage-point shift in coding assistant usage share between January and June 2026.

GenAI SaaS: Slower Growth, Fewer Apps, More Movement

Employee adoption of GenAI SaaS applications has largely plateaued in 2026. After reaching 34.2% of employees in December 2025, adoption held in a narrow band of 35.1% to 35.9% between January and May before rising to 37.7% in June, the first meaningful movement in six months.

Company size continues to shape how fast that adoption happens. Employees at small companies with 100 to 1,500 employees use GenAI SaaS applications at a **substantially higher rate** (42.3% in June) than employees at companies with more than 1,500 employees (31.5% in June).

Figure 1: Percentage of employees using GenAI SaaS, by company size, January through June 2026.



Consolidation, Not Contraction Is Changing How Work Happens

The number of distinct GenAI SaaS applications used by the average enterprise fell from 50.5 in the first half of 2025 to 47.6 in the first half of 2026. One working hypothesis is that enterprises are consolidating around a smaller set of sanctioned tools rather than genuinely slowing their AI usage.

That hypothesis is reinforced by what happened to usage volume. Data movement events into and out of GenAI SaaS, uploads, downloads, and copy-paste actions, grew 80% year over year between June 2025 and June 2026.

Fewer applications are now carrying far more traffic, and in short, AI usage is exploding, while simultaneously consolidating into applications and tools users trust. Whether this is the result of certain application capabilities, security teams wrangling shadow AI issues, or user-specific biases toward specific tools (i.e. ChatGPT vs Claude), is unknown. Regardless, this data shows a significant shift in how data interacts with AI.




Figure 2: Data movement events in and out of GenAI SaaS, June 2025 vs. June 2026.

Endpoint Agentic AI: Faster Growth, More Fragmentation

While GenAI SaaS plateaus in adoption, endpoint agentic AI, including coding assistants and autonomous agent applications running directly on employee devices, continues to accelerate.

Developer adoption of AI coding assistants grew from **49.5%** in December 2025 to a peak of **60.5%** in May 2026 before easing to **55.5%** in June. Even accounting for that June pullback, endpoint agentic adoption has grown roughly **six times faster** than GenAI SaaS adoption over the same window.




Figure 3: Percentage of developers using AI coding assistants, December 2025 through June 2026.

Coding Assistants Are Reshuffling, Not Consolidating

Beneath that adoption growth, market share among coding assistants shifted considerably between January and June 2026.

Codex gained **16 percentage points** of usage share and Claude Code gained **6 percentage points**, while Cursor lost **11 percentage points**. These figures measure share of use among developers already on a coding assistant, not overall adoption, so a given tool can lose share even as total developer adoption keeps growing.




Figure 4: Share of AI coding assistant usage, January vs. June 2026.

Adoption share and usage intensity are not the same signal. Among companies with more than 100 Claude Code users, developers launch Claude Code an average of 6.45 times per machine per day, compared with 2.15 daily launches for Codex. A tool with a smaller or slower-growing share of developers can still generate the majority of an organization's agentic activity, and therefore carry the majority of its data exposure.

This divergence in usage intensity is also shaped by the pace of model and product releases on both sides. Anthropic has been shipping updated frontier models roughly every few weeks, while OpenAI has kept a similar cadence, moving Codex onto GPT-5.5 in April and previewing GPT-5.6 in late June. Each release

tends to shift developer behavior briefly, as teams test new capabilities against existing workflows.

But the two products have also diverged in distribution strategy: Codex ships bundled into ChatGPT plans at little or no marginal cost, which has driven rapid growth in raw developer counts, while Claude Code's usage sits behind paid tiers built specifically around deep, terminal-native, repository-level work.

That distinction matters for risk modeling, as a tool optimized for broad, low-friction distribution and one optimized for deep, repeated engagement with a codebase will produce very different exposure profiles even at similar headcounts, and the launch cadence of new models is likely to keep widening rather than closing that gap.

Beyond Coding: Endpoint Agents Are Broadening

Endpoint agentic activity is no longer confined to developer tools. Looking at Claude Desktop, ChatGPT Desktop, and Microsoft Copilot Desktop as a category, the number of users grew **43.2%** between January and June 2026. That growth was not evenly distributed: Claude Desktop usage grew **1,233%**, while ChatGPT Desktop declined **33.2%** and Microsoft Copilot Desktop declined **81.8%**.

The scale of Claude Desktop's growth, off a small January base, means the category is still early relative to GenAI SaaS. But the direction is clear: **autonomous, agentic applications, not just chat interfaces, are becoming the default way employees interact with AI on the endpoint.** Each of these applications can read local files, execute actions, and move data with a level of system access that browser-based GenAI SaaS tools do not have.




Figure 5: Non-coding endpoint agentic app usage, January vs. June 2026

What This Means: From Tool Sprawl to Workflow Risk

Read together, these two halves of the data point to a single conclusion: the center of gravity in enterprise AI is moving from applications to workflows. GenAI SaaS adoption is stabilizing around a smaller, more concentrated set of tools that carry more traffic than ever. Endpoint agentic AI is growing faster, fragmenting across more tools, and operating with deeper access to local data, systems, and code.

A security program built to track which GenAI SaaS applications employees use, at what rate, and with how much sensitivity, needs a second, complementary line of sight into which agentic applications run directly on the endpoint, how often, and with what level of system access. Adoption-rate metrics alone undercount risk, because usage intensity and the workflows a tool touches vary widely within the same adoption number. Two organizations with identical AI coding assistant adoption rates can carry very different risk profiles depending on which assistants their developers actually run, how often, and against which repositories.

This is the shift data security programs now need to address directly: not just which AI tools an organization sanctions, but how sensitive data actually moves through the human and agentic workflows those tools now touch.

AI Security for a Workflow-First World

Six months of new data confirm what the original 2026 AI Adoption & Risk Report projected: the center of enterprise AI risk is moving from a fixed list of sanctioned applications toward the workflows, human and agentic alike, that those applications now operate inside of. Understanding that shift requires visibility that follows data through a workflow, not just visibility into which applications are installed.

Cyberhaven's approach to data security is built around that same principle, tracing data lineage across the applications, endpoints, and agents that make up modern work. As organizations plan their AI security posture for the second half of 2026, that workflow-level view is the piece most governance programs still lack.

Learn more about Cyberhaven's approach to data security for the agentic enterprise at cyberhaven.com.

[Request a demo](#)

About Cyberhaven

Cyberhaven is the leader in Data Security for the Agentic Enterprise, trusted by the world's most innovative companies to protect their most sensitive data. Cyberhaven traces the full lifecycle of your data, adapting protection as context changes. Cyberhaven protects workflows, not just data, by connecting lineage, identity, and behavior to understand how data is created, copied, fragmented, and shared across workflows. Operating in real time where humans and agents act on data, Cyberhaven helps stop risk before it becomes loss. AI changed work. Cyberhaven protects it. Learn more at cyberhaven.com.