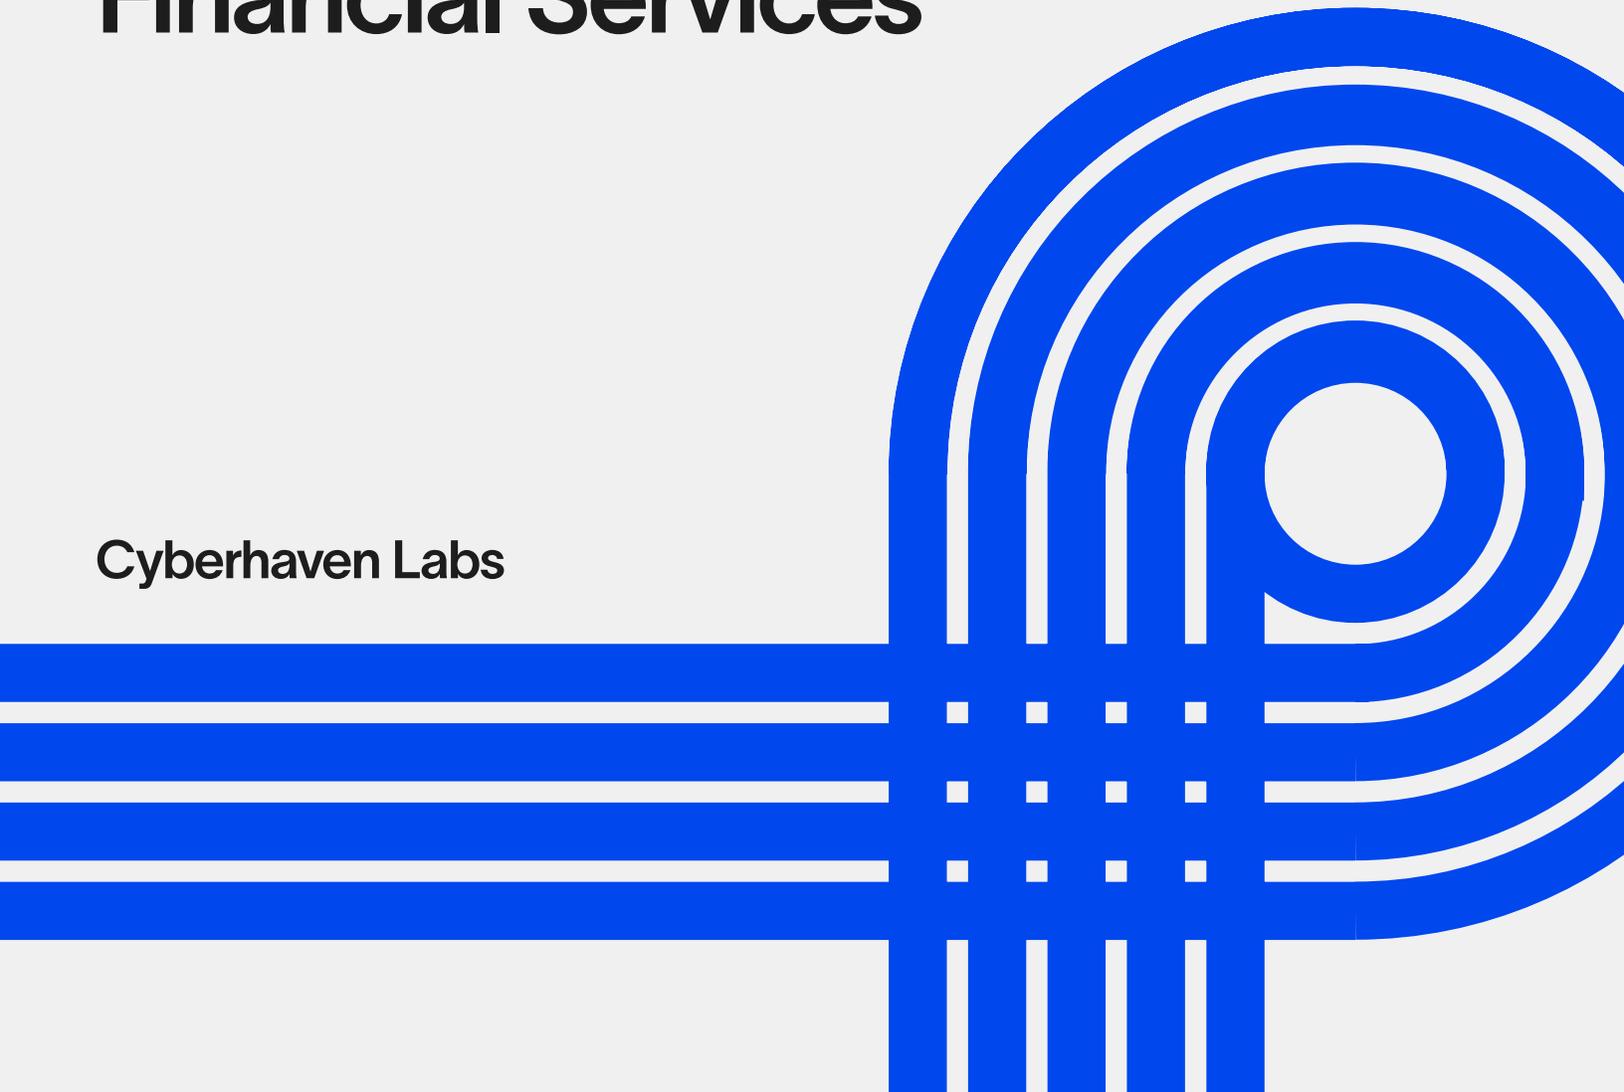




2026 AI Adoption & Risk Report

Financial Services

Cyberhaven Labs



Introduction

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

What began as individual experimentation has rapidly evolved into tools embedded in core business workflows. The speed of this shift, from novelty to operational dependency, has often outpaced enterprises' ability to understand, govern, and secure AI usage.

Cyberhaven Labs' 2025 research captures this shift, analyzing usage across generative AI SaaS, endpoint AI applications, and AI agents. Leveraging billions of real-world data movements from 222 companies, we measured adoption via active users and event-level activity, revealing not just presence, but deep integration into workflow operations.

While traditional chat-based GenAI SaaS adoption is plateauing, AI coding assistants, browser-based agents, and custom AI agents are growing rapidly, operating inside development environments and workflows with limited oversight.

As AI becomes infrastructure rather than a standalone interface, security risks intensify: employees are inputting source code, financial data, and intellectual property into a fragmented ecosystem of tools, many outside traditional IT visibility, across personal accounts and SaaS platforms without enterprise-grade security. Shadow adoption increases, controls are inconsistent, and risk accumulates faster than organizations can manage.

These risks are real, material, and concentrated among aggressive adopters.

This report provides a data-driven view of 2025 enterprise AI usage, highlighting where adoption is accelerating and where security risk is compounding across financial services enterprises. By analyzing real-world patterns across industries, departments, tools, and data types, Cyberhaven Labs helps security and technology leaders understand not just the scale of adoption, but the context needed to govern AI safely in 2026.

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 organizations¹ are experimenting with AI agents, enterprises are spending four times more on AI software than on traditional software, and 74% of executives stated² they achieve returns within the first year of AI tool deployment.

For financial services organizations, this number is even higher. A 2025 survey conducted by Google Cloud and National Research Group confirms that GenAI initiatives are still delivering, with 77% of financial services executives stating their organizations achieved positive ROI within the first year.

However, AI adoption 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 in the 99th percentile of GenAI adoption — are interacting with hundreds of GenAI applications over the course of 2025. In the most advanced cases, organizations are using more than 170 GenAI tools. By contrast, cautious enterprises, making up the 5th percentile, typically employ around 10 GenAI tools.

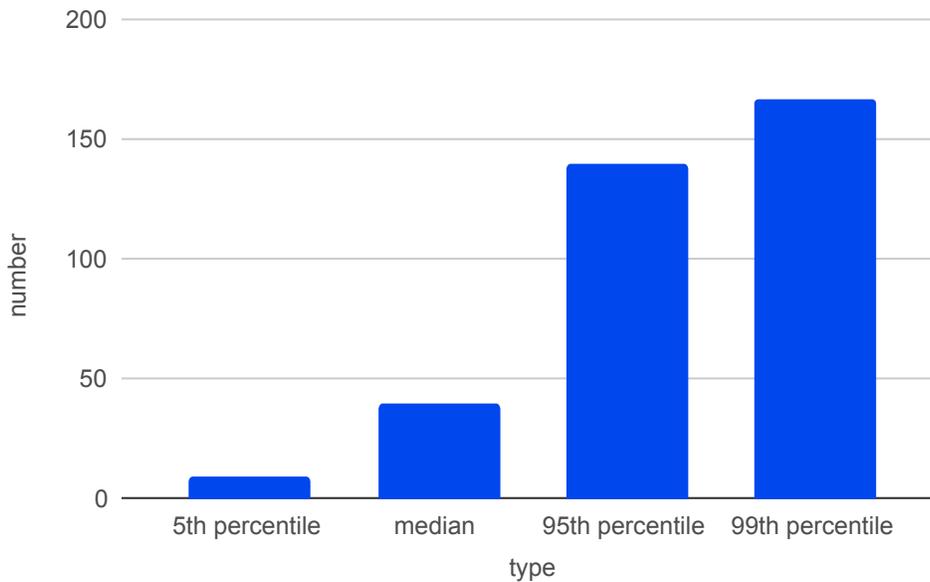
That is a 17x difference.

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

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

Generative AI SaaS

Distribution of Number of GenAI SaaS Applications Used by an Enterprise



Most Financial Services Organizations Remain Hesitant to Adopt AI

While frontier organizations are rapidly experimenting with GenAI, the majority of enterprises remain cautious. Considering the industry's risk-averse nature and the compliance and legal risks associated with adopting AI without rigorous governance, this is not a surprise.

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

Top GenAI Apps Used By Financial Services Organizations

-  chatgpt.com
-  gemini.google.com
-  claude.ai
-  perplexity.ai
-  copilot.microsoft.com
-  lovable.dev
-  app.glean.com
-  grok.com
-  notebooklm.google.com
-  chat.deepseek.com
-  otter.ai
-  ai.azure.com
-  poe.com
-  aistudio.google.com
-  app.blueflame.ai
-  chat.chaton.ai
-  web.descript.com
-  zerogpt.com
-  midjourney.com
-  elevenlabs.io
-  granola.ai
-  app.harvey.ai
-  app.synthesia.io
-  v0.dev
-  writer.com

When we compare this top 25 list to the top 25 GenAI applications used across industries, a few outliers emerge.

Financial services is the only industry to use (or at least use enough to appear in this top 25 list):

1.  web.descript.com
2.  chat.chaton.ai
3.  zerogpt.com
4.  elevenlabs.io
5.  granola.ai
6.  app.synthesia.io
7.  v0.dev
8.  writer.com

These applications favor enterprise writing, content workflows, and creative media. Missing from this top 25 list are a few of the infrastructure and development applications present in other industries such as app.devin.ai, app.codesignal.com, and huggingface.co. This aligns with the financial services industry which is text-heavy, regulated, and client-facing, meaning writing and media tools become more useful. Financial services run on documents, not code.

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, **81%** are classified as “medium,” “high,” or “critical” risk. Even when excluding “medium” risk and considering only “high” and “critical,” **62%** of tools still fall into these categories.

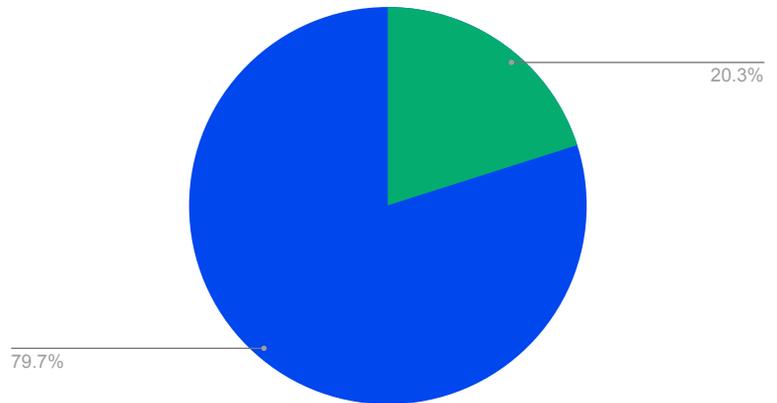
For security leaders, this means that most AI usage today occurs in tools that would not meet traditional enterprise risk standards, yet employees continue to input sensitive data into them at high rates.

Agentic AI Is Seeing Rapid Adoption

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. Instead, 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.

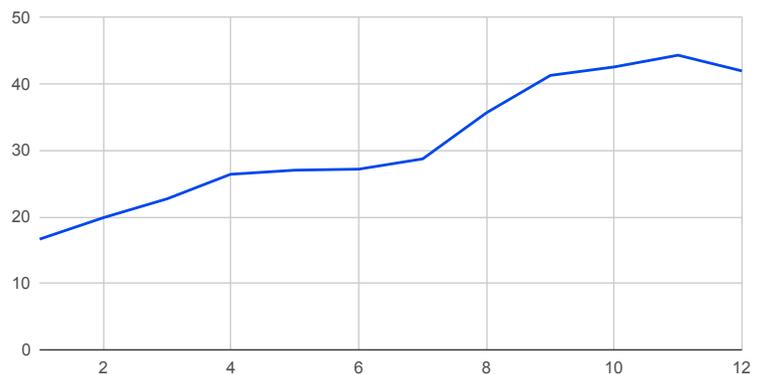
Percentage of Companies using Agent-building SaaS Platforms



If we look at the chart above, 79.7% of financial services organizations are building with agentic SaaS platforms such as Glean or CoPilot Studio. Additionally, the use of AI coding assistants (e.g. Cursor, GitHub Copilot, and Claude Code) has risen steadily for these enterprises over 2025, showing a profound interest in streamlining work functions through agentic AI.

The usage jumped from **16.7%** in January 2025 up to **42.1%** in December 2025. Put another way, the share of users within a given financial services enterprise using an AI coding agent grew by roughly **2%** each month.

Percentage of Developers Using AI Coding Assistants



Conclusion:

AI Security Is Paramount As Organizations Race To Adopt New Technology

AI adoption is accelerating, but unevenly. A small group of frontier organizations is driving rapid usage, while others move more cautiously. In many cases, innovation and experimentation are prioritized ahead of security and governance. At the same time, employees are actively using high-risk AI tools and inputting sensitive data across a growing ecosystem of GenAI applications, coding assistants, and custom agents.

For many organizations, this resembles a “wild west” environment. Tools are spreading faster than policies, usage often exceeds visibility, and sensitive data moves across systems with limited centralized control. Without clear insight into how AI is used across teams, workflows, and data types, the gap between innovation and security will continue to widen.

As adoption diverges across industries, teams, and users, AI security must become a core priority. One-size-fits-all policies are unlikely to work. Effective governance depends on understanding real usage patterns and applying controls based on data sensitivity, user maturity, and tool risk.

The challenge is complex. The scale and speed of AI adoption make manual oversight insufficient. Many organizations will benefit from specialized security solutions that unify visibility, context, and enforcement across data, users, and AI systems. Those that invest early in comprehensive AI security will be best positioned to innovate with confidence while maintaining trust, compliance, and resilience.