Nightfall AI

2026 AI Agent Risk & Action Report

AI Adoption, Innovation, and the Emerging Risk Landscape

Nightfall AI

Table of Contents

Executive Summary

Artificial intelligence has crossed a structural threshold inside the enterprise. What began as isolated experimentation with GenAI tools has evolved into something far more consequential: AI is now embedded across workflows, connected to core business systems, and increasingly acting autonomously through agents.

This report draws on anonymized, aggregated data from Nightfall's customer base across technology, financial services, healthcare, and other industries, encompassing more than 35,000 distinct applications. The data provides a point-in-time view into how AI is actually being adopted and used inside real enterprise environments — revealing patterns that are often invisible to traditional security approaches.

The core finding of this report is that AI risk is no longer defined by the tools themselves, but by the data connectivity they enable. Every integration between an AI system and a SaaS platform, every agent connected to internal systems, and every emerging protocol such as MCP introduces new pathways for sensitive data to flow. These pathways are dynamic, compound, and increasingly autonomous, making traditional approaches to data security insufficient.

Legacy data loss prevention (DLP) technologies were not designed for this model. Built for an earlier era of email and file transfers, they rely on static rules, generate high volumes of low-quality alerts, and lack visibility into the modern AI ecosystem.

Security leaders now face a different challenge. The question is no longer whether AI is being used, but how deeply it is embedded, what systems it can reach, and what data it can access. Managing this risk requires a shift from tool-centric governance to a data-centric model.

AI is now part of how organizations operate. Risk has followed the same path. The enterprises that succeed in this next phase will not be the ones that slow adoption, but the ones that build the controls to support it. In an AI-driven world, securing data—not just tools—is the only strategy that scales.

Key Findings

01 Nightfall discovered more than **35,000** distinct applications in use across its customer base, including **519 GenAI tools** and **103 AI Agent tools.**

02 The median organization runs **182 applications** and **11 AI tools,** while the largest operates over **8,000 applications** with more than **170 AI tools.**

03 Nearly half of all organizations **(49%)** are actively using or experimenting with AI agents — autonomous systems that **take action, not just assist.**

04 MCP servers grew **100× in the last 16 months** — from ~100 at launch to **10,850+,** each one **a new pathway** to enterprise data.

05 **81%** of GenAI tool usage occurs **outside the top three providers,** meaning the majority of risk accumulates in **the long tail** of lesser-known tools.

01

AI Is Now Core Infrastructure

AI has reached a tipping point. It has moved beyond experimentation into daily operations.

35,000+
Apps Discovered

519
GenAI Tools

103
Agent Tools

AI has reached a tipping point. It has moved beyond experimentation into daily operations. Organizations are no longer using a single AI tool — they are operating across dozens of AI-enabled systems embedded throughout their workflows. The median organization in our dataset uses 182 total applications and 11 AI tools. At the top end, organizations operate over 8,000 applications with 170+ AI tools.

98%
Organizations using GenAI

49%
Organizations using AI Agents

48%
Organizations using both

> *AI is no longer a tool. It is becoming part of the operational fabric of modern organizations.*

ACTION PLAN

- **Elevate AI governance to a core security priority,** not an extension of IT policy. The scale of exposure demands program-level ownership.
- **Maintain a continuous AI asset inventory** — what tools are in use, who is using them, and what data they can access.
- **Start scanning now** — monitor where sensitive data is moving and whether it's flowing into unsanctioned AI tools. Don't wait for a classification project.
- **Map the connections between AI tools and core business systems.** AI risk isn't about the tools — it's about what they can reach.

02

The AI Maturity Curve

AI adoption is not evenly distributed. Organizations fall into distinct maturity segments.

11

AI tools — median organization

170+

AI tools — largest organization

AI adoption is not evenly distributed. Organizations fall into distinct maturity segments, each with different risk profiles:

Segment	AI Tools	% of Orgs	Description
Early Explorers (P25)	1-3	~25%	Testing individual tools. Limited organizational adoption.
Scaling Teams (Median)	10-11	~50%	Adoption across teams. Emerging workflows. Mix of sanctioned and unsanctioned tools.
High-Growth Innovators (P90)	52+	Top 10%	Rapid experimentation. Multiple models and agents. AI embedded in day-to-day work.
AI-Native Leaders (P95-P99)	73-175	Top 5%	Deep workflow integration. Custom automation and agents. AI is infrastructure.

The gap between the median and the top is enormous. AI-Native Leaders operate 15-17x more AI tools than the median organization.

ACTION PLAN

- **Assess where your organization sits on this curve.** Your security controls need to match your actual maturity tier, not your aspirational one.
- **Early Explorers:** Start with visibility — discover what AI tools are already in use and what data they can access.
- **Scaling Teams:** Focus on shadow AI detection. This is the stage where unsanctioned tool use spreads fastest.
- **High-Growth and AI-Native:** Shift from reactive alerting to behavioral risk intelligence.
- **Revisit your maturity tier assessment quarterly.** Organizations are climbing this curve faster than annual review cycles can track.

03

AI Scales with Organizational Complexity

AI adoption is strongly correlated with total application footprint.

AI adoption is strongly correlated with total application footprint. Organizations with more SaaS tools adopt more AI tools, experiment more broadly, and integrate more deeply.

Average AI Tools by Application Footprint

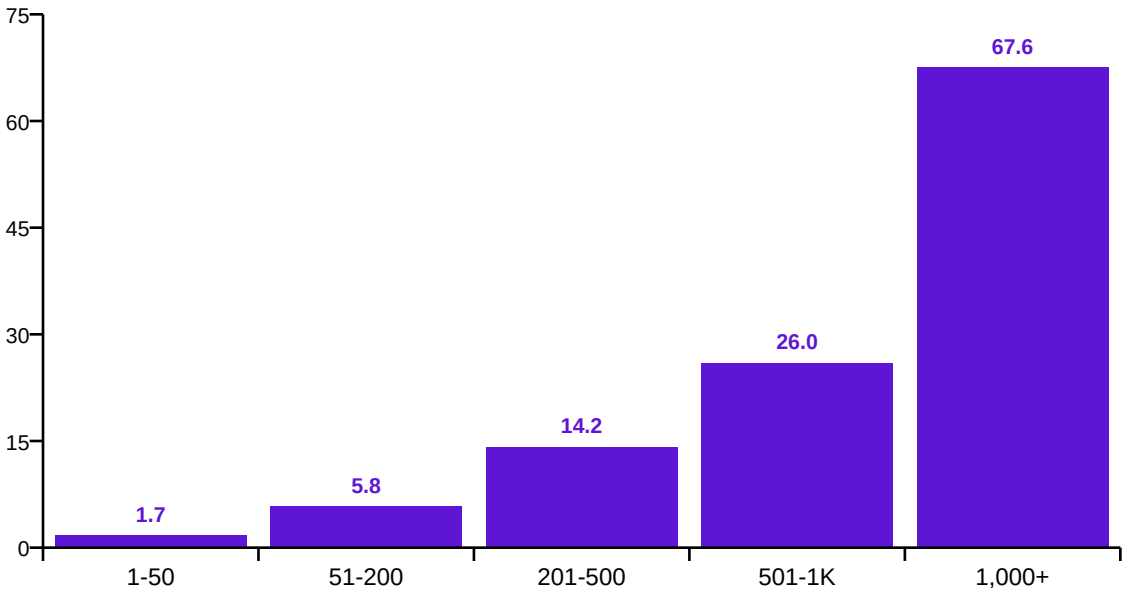

Total App Footprint	Avg AI Tools	Avg GenAI	Avg Agents
1-50 apps	1.7	1.6	0.1
51-200 apps	5.8	5.3	0.5

Total App Footprint	Avg AI Tools	Avg GenAI	Avg Agents
201-500 apps	14.2	12.8	1.4
501-1,000 apps	26.0	24.4	1.6
1,000+ apps	67.6	58.7	8.9

Organizations with 1,000+ applications use 40x more AI tools than those with fewer than 50. Agent adoption follows the same pattern but lags behind: the largest organizations average nearly 9 agents, while smaller organizations have barely begun experimenting.

> *AI is compounding on top of existing software ecosystems — not replacing them.*

ACTION PLAN

- **Treat your SaaS footprint as your AI risk surface.** Application sprawl is a leading indicator of AI exposure.
- **Add AI exposure to your app onboarding checklist.** Every new SaaS tool is a potential AI integration point.
- **Deploy API-based data protection that scales.** Per-app configuration will never keep pace with adoption.
- **If you're in the 1,000+ app tier, continuous monitoring isn't optional.** You face 40x the AI tool exposure.

04

GenAI vs. AI Agents: Two Layers of AI Adoption

AI adoption follows a layered model. GenAI tools drive productivity. Agents drive execution.

49%

Using or experimenting with AI agents

AI adoption follows a layered model. GenAI tools drive productivity. Agents drive execution.

	GenAI	AI Agents
Model	Prompt-driven	Workflow-driven
Interaction	User-initiated	System-integrated
Purpose	Assist with work	Execute work
Adoption	98% of organizations	49% of organizations
Median tools	10 per org	4.6 per org (among adopters)

Agents are not replacing GenAI — they are built on top of it. Of organizations using GenAI, 51% also use agents.

The top GenAI tools reflect consolidation around major providers:

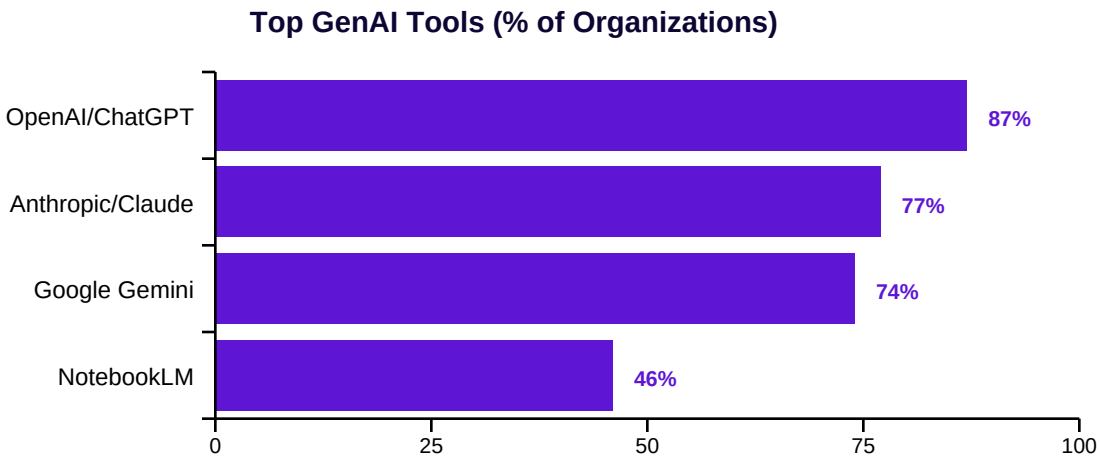

The top AI Agent tools reflect fragmentation — no dominant platform:

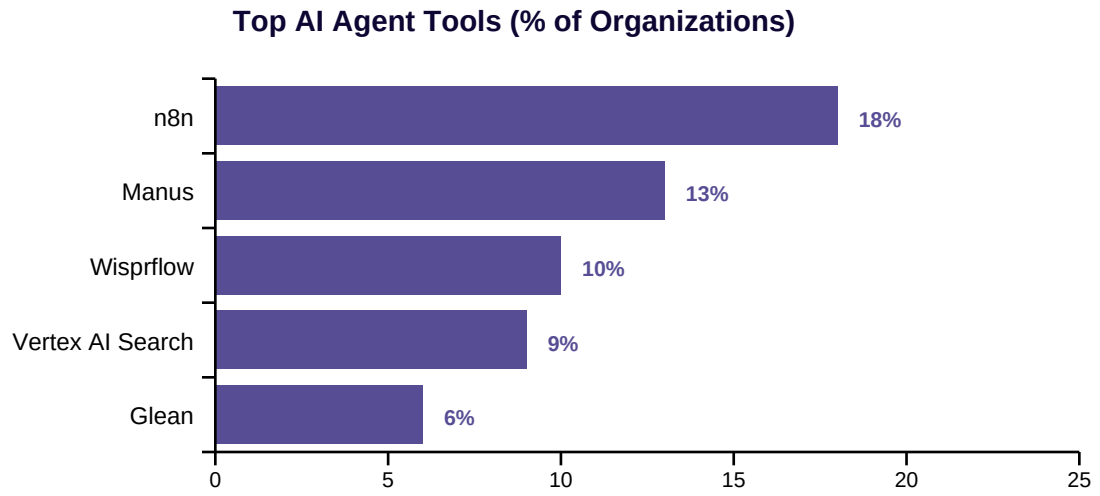

> *AI is evolving from tools that assist work to systems that execute work.*

ACTION PLAN

- **Govern GenAI and agents separately.** They have different risk profiles and need different controls.
- **For GenAI:** Deploy real-time scanning at the browser and endpoint to catch sensitive data before it enters a prompt.
- **For agents:** Get real-time visibility into what each agent can reach and treat every connection like a privileged service account.
- **Enable granular control** on agent access to systems containing regulated or sensitive data.

05

The Long Tail Problem

While a handful of AI tools dominate headlines, the real complexity lives in the long tail.

81%

Of GenAI usage outside top 3 providers

While a handful of AI tools dominate headlines, the real complexity lives in the long tail.

The AI market is splitting in two directions. GenAI is consolidating — three providers (OpenAI, Anthropic, and Google) are present in the vast majority of organizations and are becoming standardized infrastructure. AI Agents remain fragmented — no single platform has more than 18% market penetration, and with 103 distinct agent tools across our dataset, the market is early and experimental.

Yet despite consolidation at the top, the long tail tells a different story:

- 519 distinct GenAI tools observed across our customer base
- 81% of GenAI tool adoption occurs outside the top 3 providers
- The top 3 GenAI tools account for less than 20% of total adoption

The Long Tail: AI Tool Distribution

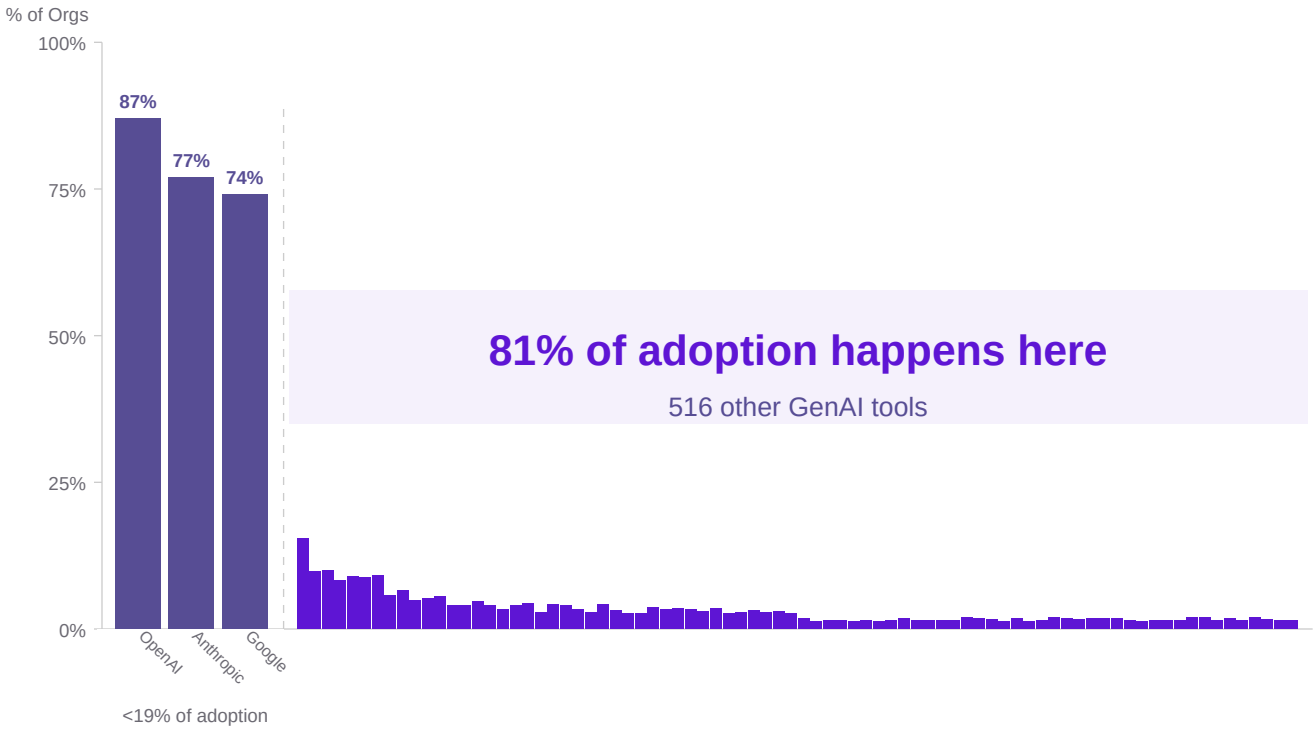

For every OpenAI or Anthropic deployment that security teams are aware of, there are dozens of niche AI tools — translation, content generation, code assistance, design, research — operating with limited visibility. The median organization uses 10 GenAI tools, but across our dataset we observe 519. The vast majority are used by only a handful of organizations, making them harder to detect, evaluate, and govern.

Top 25 GenAI Tools

#	Tool	Adoption
1	OpenAI / ChatGPT	86.9%
2	Google Gemini / NotebookLM	76.6%
3	Anthropic / Claude	74.1%
4	Perplexity	36.4%
5	Grok	27.1%
6	DeepL	25.2%
7	Gamma	24.3%
8	Microsoft Copilot	18.7%
9	ChatOn	16.8%
10	ElevenLabs	16.8%
11	GPTZero	15.0%
12	Quillbot	15.0%
13	Genspark	13.1%
14	Use	13.1%

Top 25 AI Agent Tools

#	Tool	Adoption
1	n8n	15.0%
2	Manus	10.3%
3	Wisprflow	7.5%
4	Vertex AI Search	6.5%
5	Glean	5.6%
6	HelloRetriever	4.7%
7	Make	3.7%
8	Read	3.7%
9	Sardine	2.8%
10	TrySparrow	2.8%
11	Gumloop	2.8%
12	Fireflies	2.8%
13	Devin	2.8%
14	Abnormal	2.8%

Top 25 GenAI Tools *(cont'd)*

#	Tool	Adoption
15	Adobe Firefly	12.1%
16	ZeroGPT	12.1%
17	Chatbot	11.2%
18	DeepSeek	11.2%
19	V0	10.3%
20	Higgsfield	9.3%
21	Napkin	9.3%
22	Midjourney	9.3%
23	Stitch	9.3%
24	Labs	8.4%
25	Suno	8.4%

Top 25 AI Agent Tools *(cont'd)*

#	Tool	Adoption
15	AgentSkills	1.9%
16	Sierra	1.9%
17	Rasa	1.9%
18	Ontra	1.9%
19	PhantomBuster	1.9%
20	Claude Slack	1.9%
21	Humbot	1.9%
22	GetBrick	1.9%
23	GetUnblocked	1.9%
24	Dust	1.9%
25	Copilot Studio	1.9%

Adoption rates shown as percentage of AI-adopting organizations.

> *The model layer is stabilizing, but the workflow layer is expanding rapidly.*

ACTION PLAN

- **Don’t limit the scope of your AI security to major providers.** You’ll have visibility into less than 25% of your actual AI surface.
- **Enforce policy on the data, not the tool.** A data-centric control model works regardless of destination — including tools that didn’t exist last week.
- **Assume your inventory is incomplete.** The long tail moves too fast for manual tracking. Continuous, automated discovery is the baseline.
- **Deploy browser and endpoint-level detection** to surface shadow AI usage before sensitive data reaches unvetted tools.
- **When you discover a long-tail tool, ask what data it touched first.** Blocking drives usage underground. Understanding the exposure drives a better decision.

06

The Rise of Invisible AI

AI is no longer a destination. It is embedded inside the tools organizations already use.

AI is no longer a destination. It is embedded inside the tools organizations already use:

- Productivity tools (Google Workspace Gemini: 73% of orgs)
- Collaboration platforms (Slack, Notion, Figma — each in 60%+ of orgs)
- CRM and business systems (Salesforce: 51% of orgs)
- Developer environments (Atlassian, GitHub — widely adopted)

Users no longer consciously “switch to an AI tool.” AI is invoked automatically within trusted systems. Data flows silently between applications and AI models without explicit user awareness.

This creates a new risk scenario: Invisible AI Data Exposure — where sensitive data is processed by AI within trusted systems, users may not realize AI is involved, and traditional controls lose visibility.

> *Invisible AI Data Exposure represents a fundamentally new risk category.*

ACTION PLAN

- **Audit your major SaaS platforms for embedded AI features.** Assume they are active by default unless explicitly disabled.
- **Review AI data processing terms** for your highest-risk platforms. Know where AI-processed data goes and how long it's retained.
- **Detect sensitive data moving within SaaS platforms,** not just data leaving your perimeter.
- **Make embedded AI visible to employees.** They can't flag what they don't know is running.

07

AI Risk Is Defined by Data Connectivity

This is the most important finding in this report.

This is the most important finding in this report.

AI risk is no longer about AI tools. It is about the systems connected to them.

Among organizations that use AI, 98.7% also use Business SaaS or Core System applications — the systems where sensitive data lives.

The average organization in our dataset operates:

Category	Avg Apps per Org	Orgs Using
Business SaaS	133	98%
Core Systems	26	96%
GenAI	19	98%
Developer Tools	43	89%
Cloud Infrastructure	17	92%
AI Agents	5	49%

The pathways between AI and enterprise data are growing. Every Business SaaS tool connected to an AI workflow is a potential vector for sensitive data exposure. Every Core System integrated with an agent is a potential pathway for unauthorized action.

Your AI risk surface is defined by your data layer — not your AI tools.

ACTION PLAN

- **Start with your data, not your AI tools.** Identify your most sensitive systems first, then map which AI can reach them.
- **Know the answer to 'what could this AI touch?'** For any AI tool or agent, you should be able to trace what data it has accessed or can access.
- **Tier your AI connectivity.** Define which data systems require human approval before AI can connect.
- **Build connectivity into incident response.** When an AI-related incident occurs, the first question is always 'what data was reachable?'

08

The Rise of AI Connectivity Infrastructure

The Model Context Protocol marks a structural shift in how AI systems access enterprise data.

The emergence of the Model Context Protocol (MCP) marks a structural shift in how AI systems access enterprise data. Rather than operating in isolation, modern AI systems are increasingly connected to external tools, APIs, and internal systems via standardized interfaces.

Since its introduction in late 2024, the MCP ecosystem has grown at an extraordinary pace — from approximately 100 servers at launch to over 10,000 indexed servers within 16 months.

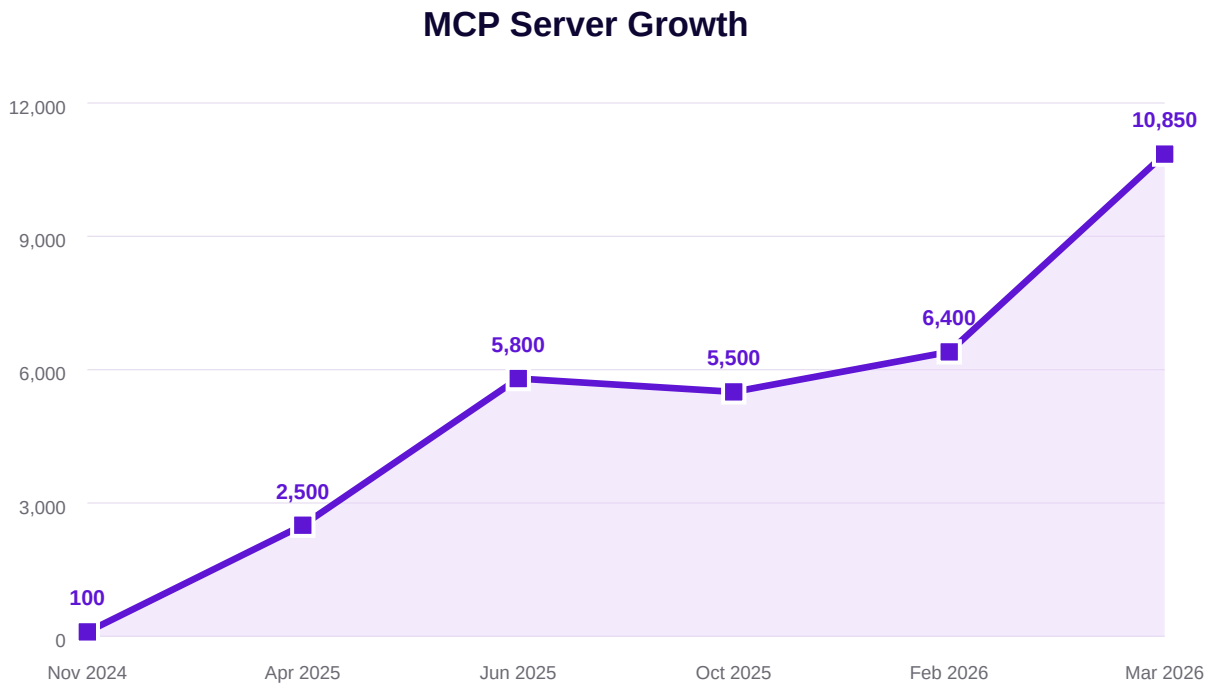

Timeframe	MCP Servers	Milestone
Nov 2024	~100	MCP launched by Anthropic
Apr 2025	~2,500	Smithery registry; OpenAI and Google add MCP support
Jun 2025	~5,800	Glama directory reaches 5,867 servers
Oct 2025	~5,500+	Official MCP Registry launched; PulseMCP at 5,500+
Feb 2026	~6,400+	Official registry at 6,400+; Smithery at 7,300+
Mar 2026	10,850+	PulseMCP / mcp.so — continued rapid expansion

This growth matters because each MCP server represents a new pathway between an AI system and enterprise data — a CRM, a database, a code repository, an internal API. The number of connections between AI and sensitive data is growing exponentially, even when the number of AI tools stays flat.

Controlling the agentic integration layer

MCP-connected agents create a new integration layer between AI tools and enterprise systems that traditional DLP rarely sees. Through MCP, AI tools can access SaaS applications, repositories, databases, and local environments without an employee manually moving the data. The data moves automatically, at machine speed, across system boundaries, without generating the user-initiated events that traditional controls rely on to detect risk.

Monitoring every step of these workflows in real time is complex and noisy. The more effective control point is the data being shared with AI tools and the systems they are connected to. Securing this layer requires visibility into connected MCP servers, what SaaS apps and systems they expose, and whether sensitive data is being shared.

> *AI connectivity infrastructure is scaling faster than AI adoption itself. The risk surface is no longer defined by how many AI tools an organization uses — it is defined by how many systems those tools can reach.*

ACTION PLAN

- **Discover MCP connections now.** Assume MCP-connected tools are already active in your developer and power-user populations.
- **Maintain an approved MCP server registry.** Evaluate new MCP connections before they reach production systems.
- **Apply the same data policies to agent-initiated flows as human-initiated ones.** The risk is identical. The volume is far higher.
- **Control the MCP gateway.** Blocking unauthorized MCP connections is the highest-leverage intervention for any organization with active agent adoption.

09

Compound AI Systems

Organizations are building compound AI systems that interact in complex ways.

Organizations are no longer using AI in isolation. They are building compound AI systems — combinations of GenAI tools, agents, embedded AI, and connected business systems that interact in complex ways.

The data makes this clear:

- 89.5% of GenAI adopters use 2+ AI providers (multi-model)
- 68.4% use 5+ distinct GenAI tools
- 51.3% layer agents on top of GenAI
- Nearly all connect AI to core data systems

This means risk is no longer linear. A single data exposure event can cascade through:

1. A GenAI tool that processes a sensitive document
2. An agent that acts on the output
3. A Business SaaS system that receives the result
4. A Core System that stores it permanently

> *AI environments are becoming interconnected systems — not individual tools. Risk assessment must account for the compound effect of AI + data + connectivity.*

ACTION PLAN

- **Model multi-step threats, not single events.** Your threat models need to account for data flowing across AI tools, agents, and business systems in sequence.
- **Require security review** for any workflow chaining multiple AI systems — especially where an agent acts on GenAI output.
- **Set data handling rules at the workflow level.** What can enter, what can exit, and what must never persist?
- **Stress-test your incident response.** Can your team reconstruct a data path across three AI systems and two SaaS platforms in under an hour?

10

AI Adoption Across Industries

AI adoption is high across all industries, but the patterns differ.

AI adoption is high across all industries in our dataset, but the patterns differ:

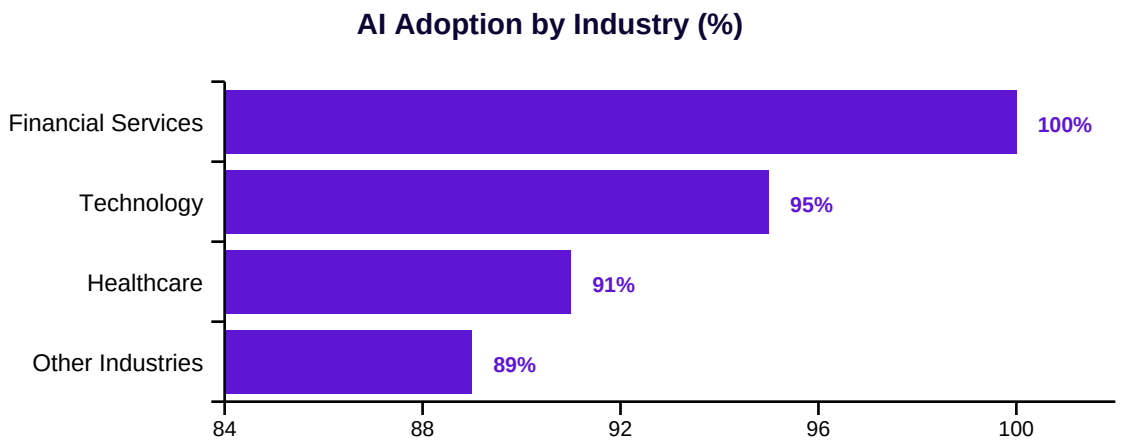

Industry	AI Adoption	Key Pattern
Financial Services	100%	Regulatory exposure. Every org has adopted AI despite compliance complexity.
Technology & Software	95%	Rapid experimentation. Highest tool counts and deepest integration.
Healthcare	91%	Data sensitivity. High adoption despite HIPAA and PHI concerns.
Other Industries	89%	Broad adoption. AI is not limited to technology-forward sectors.

By company segment:

Segment	AI Adoption
Enterprise	100%
SMB	95%
Commercial	86%

The key question is no longer whether an industry has adopted AI. It is how AI is being used — and what data it touches.

> *Every industry faces AI risk. The difference is the nature of that risk: regulatory exposure in finance, data sensitivity in healthcare, tool sprawl in technology.*

Nearly half of Nightfall's customers are in healthcare and financial services, the two most regulated data environments in the world, and our detection engine is trained to identify the sensitive data types that matter most in each vertical.

ACTION PLAN

- **Anchor your AI data policy to your highest-risk data types.** Identify the two or three that carry the greatest regulatory or business exposure.
- **Financial services:** Ensure AI-connected workflows are in PCI scope and included in your next compliance audit.
- **Healthcare:** Confirm BAA coverage extends to every AI vendor that may process PHI — including embedded AI in SaaS platforms.
- **Technology:** Treat credential and source code exposure via AI as Tier 1 risk. A single leaked API key can dwarf a typical data breach.

11

The Nightfall AI Risk Index

A composite measure to track how AI risk evolves over time.

To help organizations track how their AI risk evolves over time, we propose the Nightfall AI Risk Index — a composite measure built from the metrics in this report:

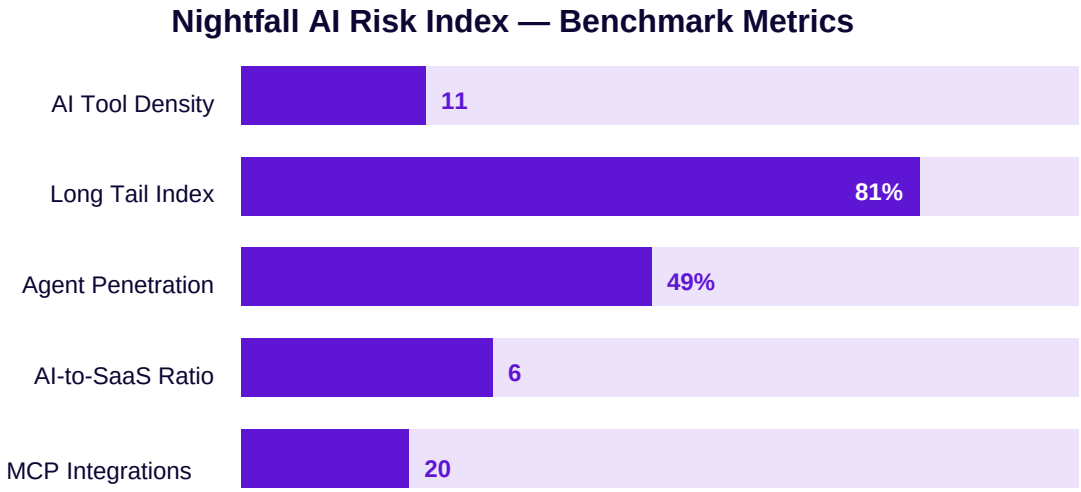

Metric	What It Measures	Benchmark
AI Tool Adoption	Active AI tools utilized across the company	11
Long Tail Index	% of AI tools outside of OpenAI, Anthropic, Google, Microsoft	81%
AI Agent Adoption	Active AI agents utilized across the company	5
AI-to-SaaS Ratio	AI apps as % of total application footprint	~6%

Metric	What It Measures	Benchmark
MCP Integrations	Active MCP servers utilized across the company	20

Future editions of this index will incorporate shadow AI detection rates, sensitive data exposure paths, and AI growth velocity as Nightfall App Intelligence expands.

Organizations scoring high across these dimensions face the greatest AI risk surface — not because AI is inherently dangerous, but because the pathways between AI and sensitive data are multiplying faster than governance can keep up.

> *The index provides a framework for tracking AI risk over time.*

ACTION PLAN

- **Score your organization against each index and establish a baseline.** The trend matters more than any single snapshot.
- **Use the index to communicate AI risk to the board.** Each metric maps to a business risk question that non-technical stakeholders can engage with.
- **Let your scores drive your priorities.**
- **Reassess quarterly.** Annual review cycles can’t keep pace with AI adoption.

Conclusion

AI is now part of how organizations operate. Risk has followed the same path.

AI adoption is no longer the key question. Nearly every organization has adopted AI. The real questions are:

- How deeply is AI embedded in day-to-day workflows?
- How many systems are connected to AI tools — and what data flows between them?
- How fast is the environment evolving — and can governance keep pace?

The organizations in our dataset are not experimenting. They are operating AI at scale — with a median of 10 GenAI tools, nearly half adopting agents, and AI connected to business systems across virtually every organization.

The central thesis of this report: AI is transforming from a tool into infrastructure. Risk is following the same path. The organizations that treat AI risk as a data connectivity problem — not an AI tool problem — will be the ones that successfully support innovation without losing control.

Legacy DLP relies on static rules and pattern matching across known channels — an approach that misses context, generates alert overload, and has no answer for Shadow AI or AI agents moving data between systems autonomously. Protecting today's environments requires AI-native detection that operates across every surface where data moves — including the ones that didn't exist last quarter.

How Nightfall Helps

Nightfall AI is the only DLP platform purpose-built on AI since 2018.

Where legacy DLP delivers 5–25% detection accuracy through static regex rules, Nightfall's AI-native detection engine operates at 95% accuracy, reducing alert fatigue and letting security teams focus on real threats rather than false positives. Where legacy tools require 6+ months to deploy and tune, Nightfall deploys in days via lightweight agents, browser extensions, and API-based connectors across the platforms where your data actually lives: Slack, Google Drive, Gmail, Microsoft 365, GitHub, Salesforce, Zendesk, Notion, and more.

As the findings in this report make clear, the next frontier is the agentic layer. Nightfall's MCP and AI Agent Security capability, currently in preview, gives security teams the visibility and control they need as AI agents become long-lived actors inside the enterprise. It provides:

- Automatic discovery of MCP servers across developer environments
- Mapping of agent access to enterprise systems
- Real-time blocking of unauthorized connections before sensitive data is exposed

Nearly half of Nightfall's customers operate in healthcare and financial services, environments where the cost of a data breach is measured not just in dollars, but in regulatory action and patient or customer trust. Our detection engine is purpose-trained on the sensitive data types that matter most in regulated industries: PII, PHI, PCI, credentials, and proprietary IP.

See how Nightfall protects your AI environment with a personalized demo.
Request a demo at nightfall.ai