


State of AI Usage Report 2026


Real-world data on how enterprises use
AI tools, how users engage in conversations,
and where risk really lies



Executive Summary


#1 Even Though AI is Everywhere, Most Users Are Casual

While nearly half of enterprise employees, on average, have used AI in the workplace, only a relatively small number do so on a regular basis. Only 18% of enterprise users use AI on a weekly basis, and that number grows consistently over the time period: 30% on a monthly basis, 40% on a quarterly basis, and 48% on a longer time frame. This means that while about half of users “touch” AI tools, less than one-fifth of users do so consistently, and most users are, in fact, casual.




#2 While AI Security Concerns Are Widespread, Risk Is Concentrated In A Small Number of Power Users

Enterprise AI activity is highly uneven. While half of employees have engaged AI tools with 12 conversations or less, the top 5% of enterprise AI users have 144 conversations or more. The same dynamic holds for prompts within conversations: 50% of users typically write only 1-2 prompts per conversation, whereas the top 5% write 18. This creates a small population of “AI power users” who drive a disproportionate share of enterprise AI usage, and as a consequence, AI risk.




#3 Although ChatGPT Is (Still) the AI King of The Enterprise, Copilot M365 Is Catching Up

ChatGPT continues to dominate enterprise AI usage. While ChatGPT is used by 36% of enterprise users, it still drive 55% of conversations in the enterprise. This means employees who use ChatGPT tend to use it more heavily than other tools. Nonetheless, Copilot M365 is quickly closing the gap, with 29% user adoption and nearly 24% of conversations. Other platforms, such as Gemini, Claude, and Deepseek, are also used in the enterprise but lag significantly behind in terms of enterprise AI adoption.




#4 For All the Discussion on AI Security, Personal AI Usage Still Accounts For Nearly Half of AI Consumption

47% of enterprise AI conversations are done through personal identities rather than corporate accounts. This means they are completely outside of the scope of enterprise AI security. Moreover, 14.39% of conversations done using a corporate identity are via personal AI licenses, meaning that the data is still being trained on. This creates major visibility, governance, compliance, and data retention blind spots.




#5 Sensitive Data Exposure Through AI Is Already Widespread

More than 6% of enterprise AI conversations contain sensitive data. ChatGPT, in particular, has a high rate of sensitive data exposure (8.38%), making it the single largest enterprise AI data exposure channel. The issue is no longer whether employees will share sensitive data with AI; it is how often, where, and through which identities it happens.



#6 AI Extensions Are The AI Consumption Channel That Most Security Teams Fail to See

AI browser extensions are rapidly spreading across enterprise environments, especially in small and mid-sized organizations where roughly 15% employees already use at least one AI extension. Nearly 75% of AI extensions are granted high or critical permission levels, and 16.31% already have known CVEs, making AI extensions disproportionately risky compared to the broader extension ecosystem.



CISO Recommendations

#1

Start With The Browser - That's Where Most AI Usage Happens

AI is no longer limited to ChatGPT tabs. Employees increasingly interact with AI through embedded copilots, SaaS applications, AI extensions, AI search tools, and connectors, most of which operate directly inside the browser. While major platforms like ChatGPT and Copilot dominate usage, the bigger visibility gap increasingly comes from the growing long tail of smaller AI tools that security teams are not monitoring. Organizations need to start AI governance where most AI activity actually happens: the browser.

#2

Identify and Monitor High-Risk AI Power Users

AI activity is highly concentrated among a small percentage of power users. Security teams should identify heavy AI users and apply adaptive monitoring and controls based on their interaction volume, data sensitivity, and platform usage patterns.


#3

Shift From "Block or Allow" to Inline AI Guardrails to Prevent Data Exposure

Blocking AI outright is no longer realistic, and an "allow-all" approach is equally risky. Organizations need inline monitoring and real-time guardrails that inspect prompts, uploads, responses, and user actions as they happen, allowing employees to use AI productively while preventing sensitive data exposure, policy violations, and risky interactions in real-time.

Enterprise AI Usage Is Mainstream, But Usage Rates Vary

Nearly 1-in-5 employees regularly interact with AI assistants on a weekly basis, while over 30% use them monthly, 40% engage with them on a quarterly basis, and nearly half of enterprise users (47%) interact with AI tools over a longer time frame. This shows that AI is no longer limited to early adopters or isolated technical teams. However, most users are not daily/weekly users but use AI more casually.





Why it matters: AI usage is now common across the enterprise, but for most employees, it’s still occasional and unstructured rather than part of their daily workflow. This creates a difficult governance problem: AI is widespread enough that companies can’t realistically block it, but because usage is fragmented and often casual, a large portion of it happens outside centralized oversight and visibility.


ChatGPT Remains the Most Used AI Tool, But Copilot M365 Is Closing the Gap

ChatGPT is the clear leader in enterprise AI usage, both in adoption and engagement. While about 36% of enterprise users interact with ChatGPT, it drives over 55% of all AI conversations. This gap between user adoption and conversation share shows that ChatGPT users are significantly more active and engaged than users of other AI platforms.

Microsoft’s enterprise-focused offering Copilot M365 follows with 29% of users and 23% of conversations, reflecting broad enterprise adoption but lower interaction intensity. Gemini and Claude show moderate adoption with lower conversation shares relative to users. DeepSeek remains extremely limited in enterprise adoption, with on 0.42% of users and 0.26% of conversations.



Moreover, Microsoft shows significantly stronger enterprise AI standardization than Google. Copilot M365 sees far higher adoption than the standalone consumer Copilot offering, with 29% of users compared to just 3% for Copilot. Gemini shows the opposite trend, where the consumer version (13%) is used far more than Gemini Enterprise (5%). This suggests Microsoft users are more likely to operate within enterprise-governed AI environments, while many Gemini users still rely on consumer AI access paths outside centralized governance.




Why it matters: Despite the constant hype around new AI platforms, enterprise AI usage remains heavily concentrated around ChatGPT and Copilot M365. Most other tools, including Gemini (including Gemini’s Enterprise version), Claude, and DeepSeek, still trail behind in terms of enterprise adoption and engagement.

Enterprise AI Usage Is Driven by a Small Number of Power Users

The average enterprise user has over 36 AI conversations, but usage is heavily uneven across employees. While the bottom 50% of users have only 12 conversations or fewer, the top 5% of users generate 144 conversations. This skewed usage distribution highlights how uneven enterprise AI usage really is and shows that a relatively small group of employees drives a disproportionate amount of enterprise AI activity.

Conversation depth is also significant: AI interactions average 5 prompts per conversation, while the top 5% of conversations contain at least 18 prompts. This is a very interesting finding that indicates that for most casual users, AI interactions are a simple one-off query (not much different than using a search engine bar), for power users they are extended multi-step interactions.



Why it matters: Enterprise AI risk is not evenly distributed across employees. Instead, organizations are seeing the emergence of “AI power users” who rely heavily on AI tools as part of their daily workflows. From a security and governance perspective, these users become disproportionately important because they also represent the highest concentration of potential sensitive data exposure and unmanaged AI interactions.

The Real Problem Isn't AI Usage, It's Sensitive Data Flowing Out Of AI

Do you know what data your employees are leaking into AI tools?

On average, enterprise users have over 36 AI conversations. However, 6.48% of all enterprise AI conversations contained sensitive data.



Why it matters: Every AI conversation containing sensitive information creates potential risk involving:


- Intellectual property leakage
- Regulatory exposure
- Data residency violations
- Customer data exposure
- Uncontrolled model training

The combination of massive AI adoption and continuous sensitive data sharing creates a new enterprise data leakage channel that needs to be handled.

Personal data appears in 5.81% of all AI conversations, making it by far the most commonly exposed sensitive data category. This includes information such as employee names, email addresses, phone numbers, customer records, and other personally identifiable information (PII).

In comparison, financial data appears in 0.96% of conversations and includes items such as payroll information, banking details, credit card numbers, and financial forecasts. IT and security-related data appear in 0.94% of conversations and include sensitive technical information such as IP addresses, SSH keys, API tokens, etc.

AI Conversations Containing Sensitive Data, by Category




Why it matters: While the rates of conversations sharing sensitive information may seem low, the sheer scale and volume of AI usage in enterprises create a broad threat surface. This creates growing privacy, compliance, and regulatory risks, especially as AI adoption expands across departments handling sensitive user and business information.

Sensitive data exposure rates differ substantially between AI platforms. DeepSeek shows the highest exposure rate, with 12.63% of conversations containing sensitive data, followed by ChatGPT (8.38%) and Copilot (8.31%). Claude (6.43%) and Gemini Enterprise (6.14%) fall in the middle, while Copilot M365 shows the lowest exposure rate at 3.65%.



Why it matters: Not all AI platforms are used in the same way or under the same governance conditions. Consumer-oriented AI tools tend to see higher sensitive data exposure rates, while enterprise-integrated platforms like Copilot M365 appear to operate within more controlled environments. This highlights the need for platform-specific AI governance, as the risk profile can vary significantly depending on where employees interact with AI.

AI Conversations by Enterprise Users Containing Sensitive Data, by Platform




ChatGPT accounts for 55.08% of all enterprise AI conversations while also showing a high sensitive data exposure rate of 8.38%, making it the single largest channel for sensitive data sharing into AI platforms. Copilot M365 represents the second-largest conversation volume at 23.61% but has a significantly lower exposure rate of 3.65%, while DeepSeek shows the highest exposure rate overall at 12.63% despite minimal adoption.




Why it matters: Risk is driven not only by exposure rates, but also by usage volume. Even platforms with moderate exposure percentages can become major enterprise data leakage channels when conversation volume is extremely high.

Conversations vs. Sensitive Conversations by Platform



Most Enterprise AI Usage Happens through Personal Accounts, Creating Invisible Governance Gaps

Almost half of all enterprise AI conversations (47.11%) happen through personal identities rather than corporate-managed accounts (52.89%). This means a large portion of enterprise AI activity may be happening outside the organization’s direct control.




On top of that, 14.39% of conversations that use corporate email identities are tied to personal AI licenses.

Such instance create a false sense of security, since employees may think that they are using a corporate account, since the login is using a corporate identity. However, it is effectively no different from using personal accounts, since organizations have no control over how data is stored, retained, governed, or potentially shared with 3rd-party providers, and that data is still used for model training.

It also highlights how frequently employees mix work with consumer AI subscriptions and unmanaged AI environments.

While the overall data shows that personal and corporate usage are split almost evenly, the breakdown across platforms shows a significant divergence. While some platforms, such as ChatGPT, Claude, Copilot, and Deepseek, are dominated by personal usage with over 60% of conversations on each tied to personal accounts, other enterprise-focused platforms, such as Gemini Enterprise and Copilot M365 are used almost entirely through corporate-managed accounts, accounting for 98.15% and 90.55% of conversations respectively.




Why it matters: AI tools are increasingly being bundled with productivity tools, and in particular Microsoft M365 (Copilot) and Google Workspace (Gemini). However, while most Copilot usage in organizations uses the enterprise-oriented Copilot M365 variant (see data in the previous chapters) and is done using corporate identities, most Gemini usage is driven by the free Gemini version (the non-Enterprise one), and is dominated by personal accounts. This creates a blind spot for organizations that is impacted by the type of tools used by employees, and the accounts they use to access them.

The Enterprise AI Problem Is No Longer One Platform, It’s Hundreds of Smaller Ones

While 70.44% of users rely on a single AI assistant, nearly 30% use multiple AI platforms, including 21.16% using two tools and over 8% using three or more. This shows that AI usage inside enterprises is becoming increasingly fragmented across platforms rather than standardized around a single approved solution.




Why it matters: Multi-platform AI usage creates growing governance complexity, as each AI assistant introduces different models, policies, data handling practices, and security controls. As employees spread workflows across multiple AI tools, organizations face expanding visibility gaps and a larger surface for sensitive data exposure.













Enterprise users interact with an average of 2.24 AI applications, while the median user uses 2 AI tools. The top 5% of users interact with 6 or more AI applications, showing that a smaller group of employees relies on a broad mix of AI tools as part of their workflows.

AI usage is no longer concentrated around a single approved assistant. As employees use multiple AI applications, organizations face growing challenges around visibility, governance, data consistency, and sensitive data exposure across an increasingly fragmented AI ecosystem.




While most AI usage is concentrated around a few tools, the data points to a growing long tail of AI tools inside the enterprise, where employees increasingly adopt niche or secondary AI applications alongside mainstream platforms. This creates a Shadow AI problem, making it difficult for organizations to maintain visibility, enforce governance policies, and control how sensitive data moves across an expanding AI ecosystem.

Top 20 Most Popular AI Apps


	ChatGPT		Google Dev		Forethought
	Copilot M365		Microsoft		Kaggle
	Gemini		Deepseek		Azure.com
	Claude		Canva		Harvey AI
	Perplexity		Lovable		Vercel
	Google		DeepAI		Wordtune
	Otter.ai		Hugging Face		

Of the top 100 most popular AI applications, ChatGPT accounted for being the most used AI tool. While most AI usage is concentrated around the top 4-5 tools, the data points to a growing long tail of AI tools inside the enterprise, where employees increasingly adopt niche or secondary AI applications alongside mainstream platforms. This creates a ‘Shadow AI’ problem, making it difficult for organizations to maintain visibility, enforce governance policies, and control how sensitive data moves across an expanding AI ecosystem.




AI Extensions Are The AI Usage Channel Enterprises Don't Know About

AI extension usage is emerging across enterprises, with small and mid-size organizations showing the highest adoption, about 15% of users run at least one AI extension. While larger organizations remain more conservative or regulated, with approximately 10% of users having them.




Nearly 75% of browser extensions request high or critical permission levels (56.4% high, 16.7% critical), while only 1.4% operate with low permissions.

Why it matters: Extensions with elevated permissions can access sensitive browser data and user activity, which means a malicious or compromised extension could easily expose sensitive information or take over user sessions.





AI extensions show a higher security risk profile. 16.31% of AI extensions have known CVEs, compared to 10.8% across all extensions. AI extensions carry a significantly higher risk profile with greater exposure to vulnerabilities and elevated permissions, increasing the risk of data exposure and misuse. The higher CVE rate in AI extensions means that they require stricter vetting, monitoring, and management before being deployed in enterprise environments.



Moreover, AI extensions are nearly 3x more likely to request cookie access than the average browser extension. Also, 41.91% of AI extensions request scripting access, compared to 15.4% of all extensions.

This means that AI extensions introduce disproportionately higher privacy, security, and data-access risks and require stricter governance, review, and ongoing monitoring in enterprise environments.



AI Connectors Are Connecting AI To New Threats

AI connectors are used by a relatively small portion of employees, with 0.16% of users interacting with at least one connector-enabled AI workflow. However, users who adopt connectors tend to connect multiple enterprise applications, averaging 4.5 connectors per user. The most common integrations connect AI platforms directly to business-critical systems such as SharePoint, Microsoft 365, Atlassian, Google Workspace, GitHub, Slack, and Figma.



Why it matters: AI connectors dramatically expand what AI platforms can access beyond simple prompt interactions. Instead of employees manually sharing data, connectors can provide AI tools with direct access to emails, documents, collaboration platforms, code repositories, internal knowledge bases, and SaaS applications. This creates a much larger and more persistent enterprise attack surface, where sensitive business data can flow continuously between enterprise systems and external AI platforms.

List of Top 10 Connectors

Application	Connector
 ChatGPT	 SharePoint
 ChatGPT	 Outlook Email
 ChatGPT	 Microsoft
 ChatGPT	 Atlassian
 ChatGPT	 Microsoft (Live)
 ChatGPT	 Google
 ChatGPT	 GitHub
 ChatGPT	 Canva
 ChatGPT	 Slack
 ChatGPT	 Figma

Key Recommendations

The data in this report shows that enterprise AI usage is expanding faster than most organizations can govern it. AI is no longer limited to a few approved tools or isolated users. Employees now interact with AI through personal accounts, embedded copilots, browser extensions, connectors, and a rapidly growing ecosystem of AI applications. To reduce exposure and regain visibility, organizations should adopt a layered approach to AI governance focused on identities, data exposure, and real-time monitoring.

#1

Stop Focusing Only on “Approved AI”; Your Biggest Risk Is the Long Tail

Most organizations are focused on governing ChatGPT or Copilot, but the data shows the real visibility gap is the growing ecosystem of secondary AI tools, embedded assistants, AI search engines, browser extensions, and connectors quietly spreading across the enterprise. The biggest AI governance failures won’t come from popular AI tools, but from the dozens of smaller AI tools security teams don’t even know employees are using.

#2

Identify and Monitor High-Risk AI Power Users

The majority of employees use AI casually. The real risk is concentrated among a very small group of heavy AI users who conduct more conversations, use more AI platforms, and engage in deeper multi-prompt interactions. Security teams should continuously identify heavy AI users and apply adaptive monitoring based on interaction volume, platform diversity, and sensitive data exposure patterns.

#3

Blocking AI Is Not a Strategy; The Future of AI Security Is Inline Guardrails

The scale of AI usage in enterprises shows that blocking AI outright is no longer practical, but also demonstrates that unrestricted AI usage creates equally serious risk. The organizations that succeed will move away from binary “block vs allow” policies and towards inline AI governance: monitoring interactions in real time, understanding context, and applying controls dynamically without disrupting productivity.

#4

Restrict Personal Account Usage and Treat Personal Licenses on Corporate Accounts as Active Shadow AI


Unmanaged personal AI accounts and personal AI licenses expose sensitive enterprise workflows to uncontrolled AI environments. Enforcing corporate AI identities and blocking personal account usage helps ensure that AI interactions, prompts, and data flows remain visible, governed, and protected under enterprise security controls.

#5

Apply Targeted Security Controls to AI Extensions and Connectors

AI browser extensions and connectors can expose enterprise systems directly to external AI platforms. Organizations should continuously vet these integrations for permissions, vulnerabilities, data access scope, and behavioral changes over time.

How LayerX Can Help



LayerX is the only AI usage control platform that lets you control every prompt, agent action, and data exchange, across any application, browser, and IDE, without changing your network architecture or disrupting user experience.

It offers comprehensive visibility into all AI tools, AI embedded SaaS, AI browsers and AI application usage, regardless of who is using it, or on what device.

Moreover, LayerX goes a step further to provide end-to-end visibility into every AI conversation across all major AI platforms, capturing both prompts and responses. This gives security teams full insight into data shared with AI tools, along with the context needed to detect exposure, accelerate investigations, and ensure regulatory compliance.

It also monitors user actions such as browsing activity, login attempts, data input, and file uploads, identifying which tools are accessed, by whom, and through which accounts (corporate or personal). This enables organizations to detect unauthorized data sharing and enforce policies to prevent data leakage from authorized AI tools and other ‘shadow AI’ applications.

Moreover, LayerX permits organizations to directly detect and enforce policies on these apps in the last-mile. Organizations can define policies based on user identity, device status, website category, data sensitivity, etc., to create tailored security policies with a range of enforcement options ranging from monitoring only, to warning users with customizable messages, to masking sensitive data, to completely blocking their actions.

To learn more about how LayerX can help you manage and secure AI usage across your organization, go to www.layerxsecurity.com and schedule a demo today!