Layer

Enterprise GenAI Security Report 2025

Real-life data on how enterprise users consume GenAI tools, who uses them, and where the security blind spots of ‘shadow AI’

THE ONLY REPORT THAT ANALYZES GenAI THREATS AT THE USER'S POINT OF RISK IN THE BROWSER!

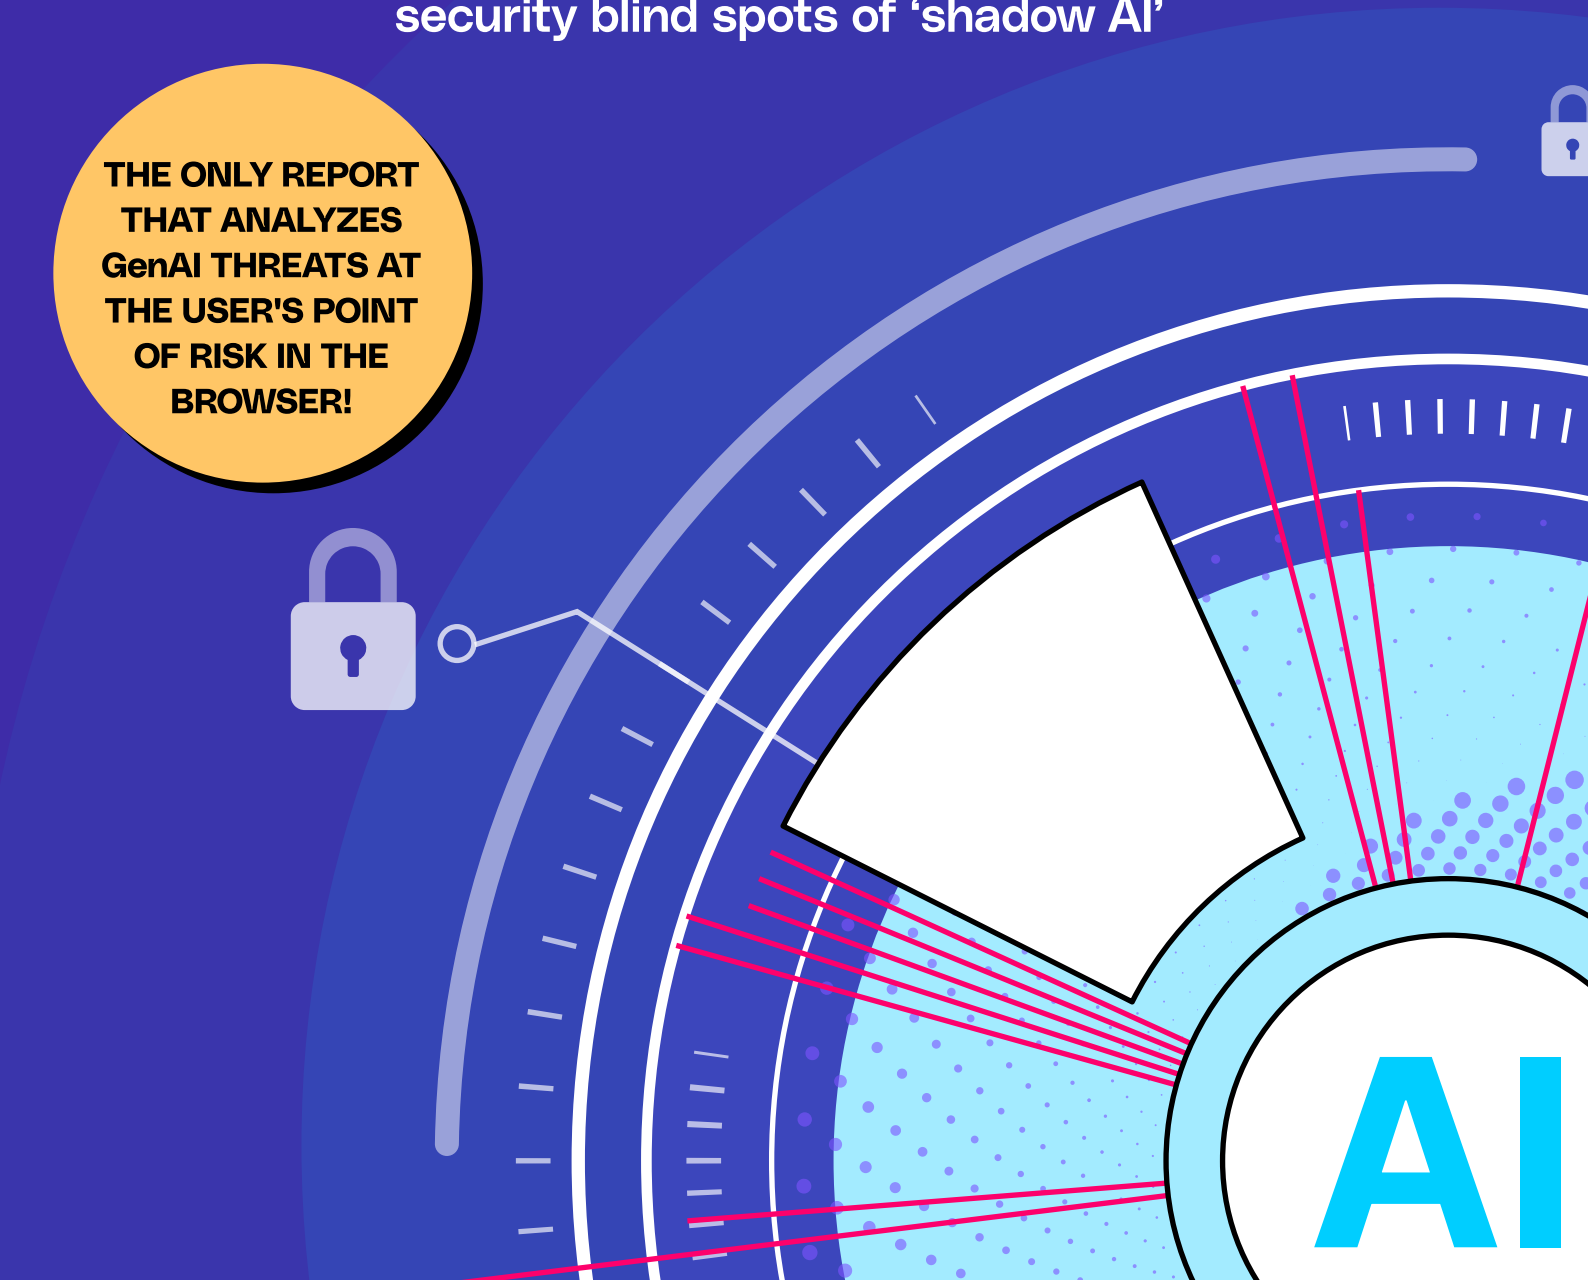

Introduction

AI has taken the world by storm.

Or has it?

Because once you move past the big, bombastic headlines and the marketing spin, there is surprisingly little hard data on how AI is actually used, especially by enterprise users.

This research is meant to do exactly that. It provides tangible statistics on how enterprise users consume AI in the workplace based on real-life data and telemetry collected from LayerX Security's customer base.

What's In This Report

This report covers several areas relevant to GenAI and AI application usage, including:

- How GenAI is used in organizations
- Who uses GenAI and AI SaaS applications in the organization
- How employees connect to and access GenAI tools
- How corporate data is shared with LLMs
- The security risks posed by GenAI browser extensions

All the findings are based on telemetry collected from LayerX's unique data set.

What Makes LayerX's Data Unique

There is no shortage of reports and surveys in the market, but what makes LayerX's data unique is where we collect our data and who we collect it from.

The LayerX Security solution is deployed directly within users' web browsers, meaning that LayerX has full visibility to all user activity and data that passes through the browser. This allows us comprehensive insights on the usage of GenAI tools and AI-enabled SaaS applications.

Moreover, LayerX's customer base is comprised almost entirely of medium and large enterprises, meaning that the insights we collect are specific to enterprises and enterprise users.

Executive Summary

How can you protect against what you don't know about?

#1

Hidden Access to GenAI Tools

Nearly 90% of logins to AI SaaS applications are done with either personal accounts, or corporate accounts not backed by SSO. Such logins don't go through organizational identity and access management systems, leaving organizations blind to their existence. Moreover, any connection to AI tools via a personal account will not be subject to organizational privacy and data controls by the LLM tool.

#2

The Long Tail of 'Shadow' AI

The top AI tools dominate over 90% of AI application usage, but once you move past the handful of best-known AI tools, there is a long tail of little-known and invisible 'shadow' AI tools that fly under the radar. Most organizations do not have visibility as to which tools are used in their organizations, by whom, or where they need to place controls.

#3

AI Browser Extensions are a 'Side Door' for Data Leakage

While many organizations already deploy (or are at least considering) dedicated AI security solutions, AI-enabled browser extensions often represent an overlooked 'side door' through which data can leak to GenAI tools without going through inspected web channels, and without the organization being aware of this data transfer.

A CISO's Framework to Security GenAI Tools

Based on these findings, we suggest CISOs and security managers implement a number of high level recommendations to cover their bases:

- **Audit All GenAI Activity:**
  Since so much of employees' AI activity is hidden, it's crucial for the organization to audit all AI activity at the endpoint level, to make sure they have visibility to it all.
- **Proactively Educate Employees:**
  Since AI is a new technology, many users are still oblivious to its associated data risks. This is why it's critical to proactively educate users and alert them to potential AI risks, as they are taking place.
- **Apply Risk-Based Restrictions:**
  While some organizations try to outright ban all AI usage, this is not a long-term solution in a world that is becoming increasingly AI-driven. This is why it's critical to apply security restrictions that are adaptive and contextual, to enable employees to use AI securely, without sacrificing productivity.

Key Findings

#1

Despite organizational security policies, organizations have no visibility into 89% of AI usage in the organization. Over 70% of connections to GenAI tools such as ChatGPT are done with users' personal accounts, even on enterprise devices. Even among logins using corporate accounts, 58% of connections are done without SSO. This means that nearly 90% of logins to GenAI tools are invisible to organizational identity access and control systems, and security and IT teams have no idea who is using GenAI tools and what data is being exposed inside GenAI conversations.

#2

Most GenAI users are casual, and may not be fully aware of the risks of GenAI data exposure. Only about 15% of enterprise employees use it on a weekly basis, and while a small percentage of users who use it extensively, most users are casual users. While some readers might see this statistic as an indication that there is no problem, we see it as a gaping hole through which users may inadvertently leak data. Exactly because most users use it casually, organizations need to ensure that their users are educated and aware of the risks.

#3

Browser extensions are the hidden threat of GenAI data leakage. About 20% of users have a GenAI browser extension installed on their computer, which can bypass AI access filters on network solutions such as Secure Web Gateways (SWGs), thereby allowing exposure of data of organizational data to remote LLMs without the organization knowing or being able to track it.

#4

Over 90% AI usage is concentrated in large, well-known apps, but there is a long tail of 'shadow' AI applications. ChatGPT alone accounts for over 50% of enterprise usage, and the top 5 AI SaaS apps for over 85% of AI usage. However, outside of the handful of well-known apps there is a long tail of lesser-used AI tools that fly under the radar. As a result, security manages don't know which other AI apps are used, and where to put controls.

#5

A small number of users expose large volumes of data. While text input is the standard form of interaction with GenAI tools, copy/paste and file upload are the channels through which data can leak at scale. Approximately 18% of users paste data to GenAI tools, and about 50% of that is company information.

GenAI Usage is Widespread, But Still Mostly Casual

14.5%

Of enterprise users access GenAI tools on a weekly basis

77%

Of user access to online LLM tools is to ChatGPT

39%

Of enterprise users who use GenAI tools regularly are software developers

The Finding

We began our analysis by looking at the top Generative AI tools such as ChatGPT, Gemini, Copilot and others. The data shows that approximately 14.5%, or about one out of seven users, use these GenAI tools on a weekly basis.

Looking at the most-used GenAI tools, OpenAI's ChatGPT is the undisputed champion, with 77% of activity, far ahead of Google's Gemini at 11%. After that, there is a minor surprise with Anthropic's Claude AI engine coming ahead of Microsoft Copilot, with 5% and 4%, respectively. Other LLMs make up the rest, with about 3% combined.

With regards to usage patterns, the data shows a wide disparity between heavy and casual users. Whereas the top 5% of 'heaviest' GenAI users access GenAI tools, on average, more than 4 times a day, the bottom 50% of users access them only 1-2 times a month. This finding indicates that while GenAI has made large inroads within a short time, most users are still casual, occasional users and that AI usage is not (yet?) a part of their day-to-day usage.

Software developers are the largest constituency of active users. Among enterprise users, 39% of users who use GenAI tools belong to R&D, 28% belong to Sales and Marketing. IT, HR, and Finance users make up single digits only. This finding is consistent with market trends of AI uses (and available tools) for software development and marketers.

Top GenAI tools by user connections

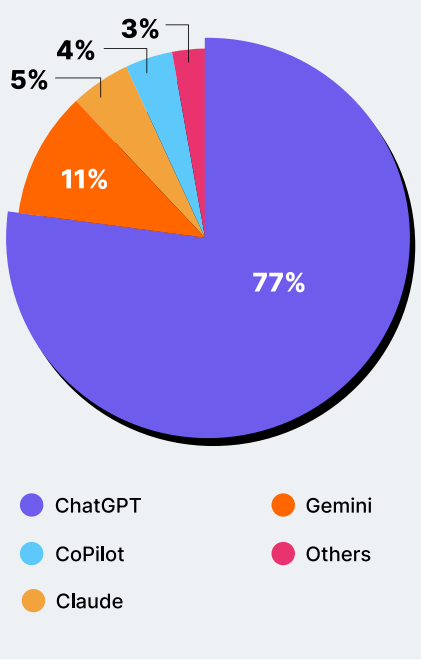

Enterprise GenAI users by department

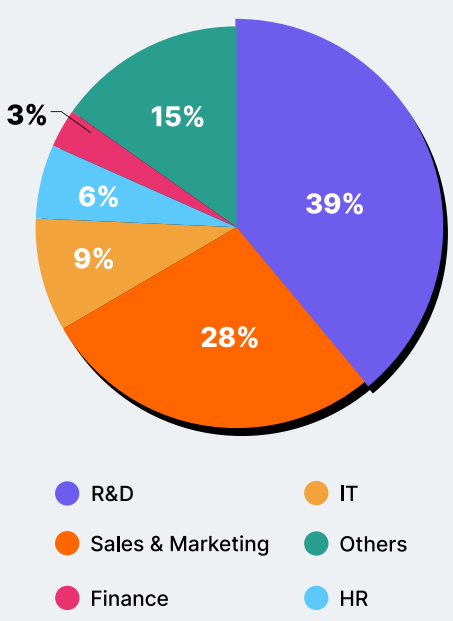

Analysis

The findings indicate that while GenAI usage is widespread, it is not (yet) ubiquitous in organizations. Nonetheless, considering the short period since ChatGPT first came into our lives in late 2022, these are significant inroads that are only expected to increase in the coming years.

Evidence of the disparity in usage between casual and heavy users could be seen in the usage patterns where the bottom 50% of users (and in all likelihood, much more than that) used GenAI tools only occasionally, whereas the top 5% of heavy users used it all the time. We expect this 'tip' of heavy users to increase with time, but for now, it shows the 'chasm' that most users have not yet crossed.

There was little surprise in the distribution of the top GenAI users of software developers, marketers and salespeople, as those are the organizational users that probably benefit the most for the 'generative' aspects of GenAI. However, as uses of GenAI expand, we expect those results to even-off.

Finally, it was no surprise that ChatGPT was the leader in terms of usage, but we were surprised to see the extent of the lead, particularly over established enterprise players such as Google and Microsoft. Whether it is first mover advantage, brand recognition, or better technology – OpenAI created a lead that will be difficult to erase in the foreseeable future.

Of course, these findings may (and will) vary among organizations, depending on the makeup of their workforce and specific line of work. However, it highlights the need for organizations to track usage of GenAI among their employees to fully understand who's using it, which tools they are using, and to what end.

The Long Tail of AI SaaS Applications

51.7%
Of all AI application access is to ChatGPT only

86%
Of all AI application access is to the top 5 AI apps

<1%
Of AI application access is to the bottom 50 apps

The Finding

Next, we expanded our analysis to look at not just GenAI tools and LLMs, but also at AI-enabled SaaS applications that are classified as 'Ai applications.'

In terms of the most commonly used AI applications, ChatGPT is far-and-away the most commonly used AI application. Among AI tool users, ChatGPT was accessed, on average, more than once per day.

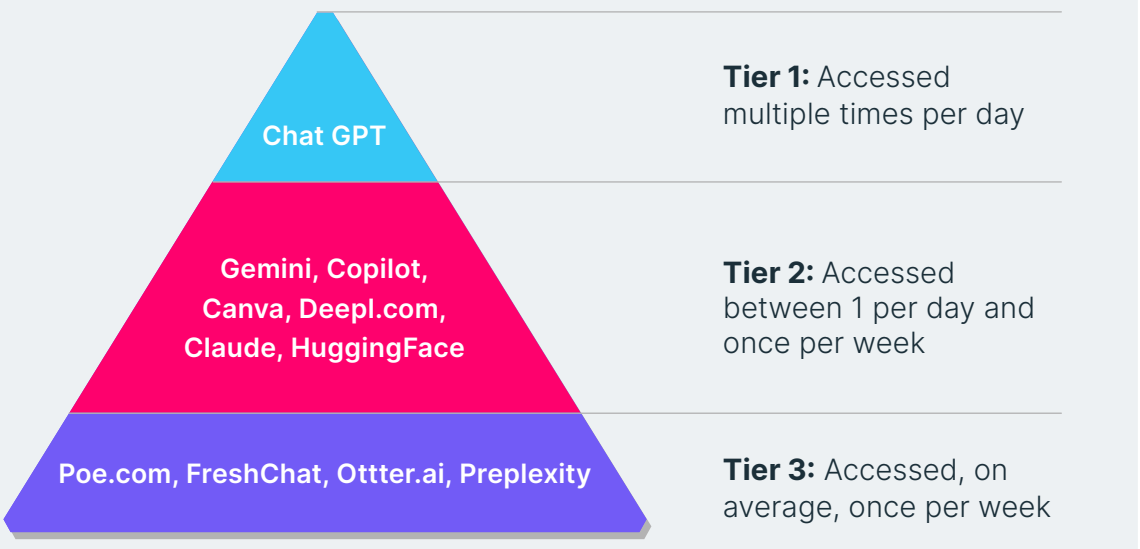

Top AI SaaS Applications by User Access

The next tier includes the other top GenAI tools, including Gemini, Claude, and Copilot, which were accessed, on average between once a day and once a week. Other AI applications in this tier include Deepl.com, Canva, and AI application marketplace HuggingFace.

The tier after that includes other well-known AI tools such as Poe.com, Otter.ai, FreshChat, and Perplexity. These tools were accessed, on average, once per week.

Of the top 100 most popular AI applications, ChatGPT accounted for 51.7% of all AI website requests. The top 5 applications accounted for over 86% of AI usage, but the bottom 50 accounted for less than 1% of requests.

These findings are an indication both of the immense popularity of ChatGPT, as well as of the 'long-tail' of AI applications, which extends beyond the top four or five AI tools that are top-of-mind to most consumers.

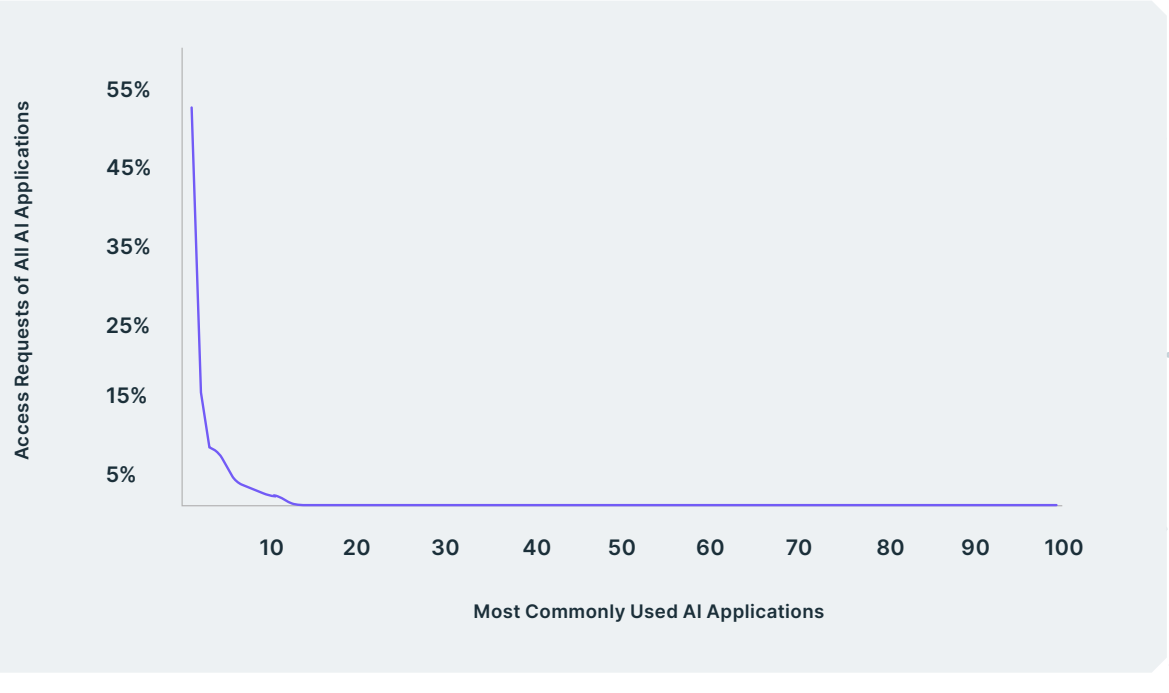

Analysis

Everybody knows ChatGPT. And the next five or ten AI tools.

But what happens when you get outside of the top 10 AI tools?

The findings show that there is an immense cliff between the handful of top tools that get the overwhelming majority of AI traffic and the rest of the pack. However, as more and more SaaS applications become AI-powered (and now ones pop-up), this leads to a long tail of 'shadow AI' applications that probably have few users and fly under the radar of corporate IT and security teams.

This means that for most organizations, apart from the few AI tools that jump to everybody's mind, there is little visibility or control over what AI tools are used in the organization, who's using them, and what data goes into them.

From an organizational security point of view, security managers need to make sure they have visibility into these 'shadow AI' applications and control over their usage.

Most Workplace AI Usage is Invisible to Organizations

71.6%

Of access to GenAI tools is done using non-corporate accounts

58.7%

Of access to GenAI tools using corporate accounts is done without SSO

11.7%

Of all AI application access is done using corporate accounts backed by SSO

The Finding

The data shows that the overwhelming majority of connections to GenAI tools are carried out using non-corporate logins. Over 71.6% of requests to GenAI tools were done using personal accounts.

Of the 28.4% of logins done using corporate accounts, 58.7% (and 16.6% of total logins) were not backed by Single Sign On (SSO). This means that only 11.7% of all logins to GenAI applications adhered to the gold standard of using a corporate account backed by SSO. The mirror image of that finding is that nearly 90% of web-based GenAI usage is invisible to the organization.

Connections to GenAI tools by account type

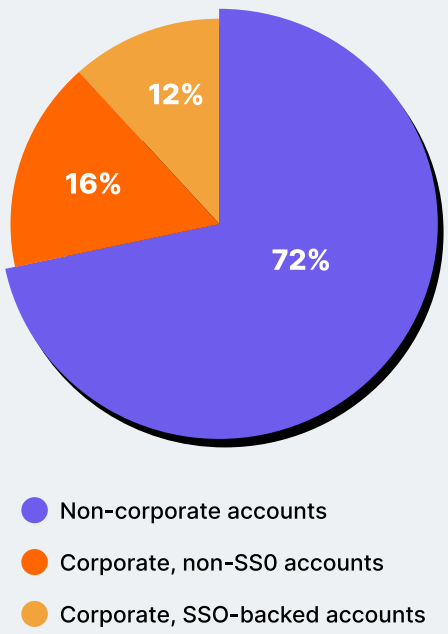

Analysis

Organizations use Single Sign On (SSO) so that corporate SaaS logins pass through the organizational IdP, giving the organization visibility into where these corporate logins are used, and providing them a measure of control (at least to the extent that they can block access to unwanted websites or SaaS applications).

However, the findings show that connections to GenAI tools by employees on organizational devices are overwhelmingly done using non-corporate (i.e., personal) accounts.

When users are connecting to GenAI tools via their personal (and typically free) accounts, they are not subject to data controls applied to corporate accounts, such as private tenants, not using data for LLM training, etc. As a result, any company information shared on such public infrastructure is compromised.

Moreover, even employees using corporate accounts do not usually use SSO. As a result, organizations have no idea of these connections. This leaves organizations blind to ‘shadow AI’ applications and the unsanctioned sharing of corporate information on AI tools.

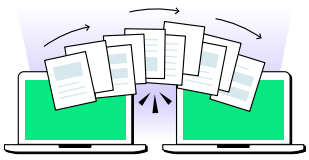

A Small Number of Users Share Large Amounts of Data

18%

Of enterprise users paste information to GenAI tools

1%

Of enterprise users upload files to GenAI tools

50%

Of paste activity to GenAI includes corporate data

The Finding

The data shows that nearly 18% of users paste data to GenAI tools, but less than 1% upload files to GenAI tools.

However, the data also indicates that users who submit data to GenAI tools via paste and file upload do so relatively frequently: among users who paste data to GenAI tools, on average, do so 6.8 times per day, and over 50% of those activities (3.8 events per day, on average) include data that could be classified as corporate information.

Although a relatively small number of users upload files to GenAI tools, those who do so are also fairly active and upload an average of 3.7 files per day.

Analysis

Text input is the standard method of interaction with GenAI tools. Virtually 100% of users input information that way into LLMs. However, that approach is typically limited in the amount of sensitive information that can be exposed since there is usually a limit to how much (and how long) a user is willing to type.

The bigger data risk, however, comes from copy/paste of information and file upload. Those are the methods in which large amounts of company information can be exposed on GenAI tools with a few keystrokes.

Approximately one in five enterprise users paste information to GenAI tools. While there is no information on where this information comes from, it makes sense that much of it is from other data sources and contains larger amounts of data (otherwise, it would have been easier to type it manually). It is no surprise, therefore, that about 50% of pasted information contained information that could be classified as corporate information.

Similarly, about 1% of users upload files to GenAI tools. While we did not review the contents of these files for this research, it makes sense that this was done for data analysis of large quantities of data, which could put this information at risk if this activity is not properly monitored and controlled.

Therefore, organizations should track user connections to GenAI tools and their activities within those tools, as well as monitor the data shared with online LLMs.

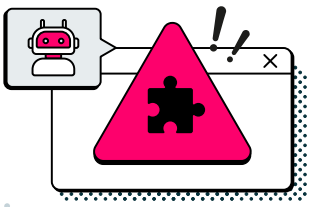

A GenAI Problem is Also a Browser Extension Problem

20%

Of enterprise users have installed a GenAI-enabled browser extensions

58%

Of GenAI browser extensions have ‘High’ or ‘Critical’ permission scope

5.6%

Of GenAI browser extensions are classified as ‘malicious’

The Finding

The research shows that 20.63% of all users have installed an AI-enabled browser extension. Of those who have such an extension installed, 45% have more than one such extension.

Of GenAI browser extensions, 58% have a permission scope classified as ‘high’ or ‘critical,’ compared to 66.6% of all extensions.

Finally, 5.6% of AI extensions are classified as ‘malicious’ and can be used to steal data.

GenAI Extensions by Permission Scope

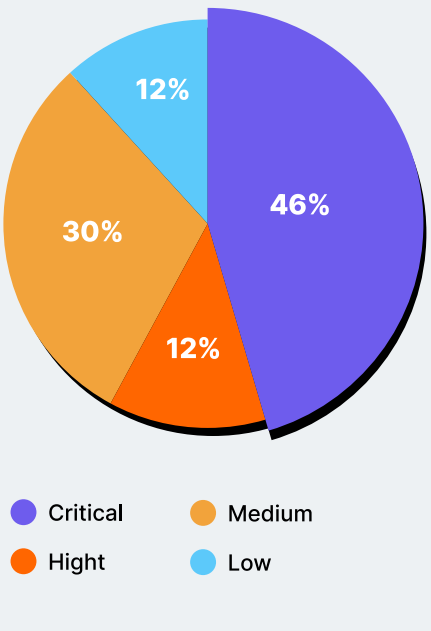

Analysis

Browser extensions are the hidden AI threat most organizations don’t know about.

While most of the attention is focused – understandably – on web access to LLMs and GenAI tools, AI-enabled browser extensions present a ‘side door’ through which data can leak out, even if the traditional web channels are blocked.

Moreover, the research shows that most AI-enabled browser extensions are granted extensive permissions to sensitive browsing information such as cookies, browsing information, web page contents, user identities, and more. This is critical since over 5% of AI-enabled browser extensions are classified as malicious.

The implication for organizations is that they need to see browser extension security as a facet of GenAI security and apply security controls over them, just as they would for web access to GenAI sites.

Recommendations

#1

Map AI Usage in the Organization

All organizations use GenAI, and most users have used GenAI tools. However, not all GenAI usage is the same: some users use it more than others, and for different purposes. This means that mapping GenAI usage in the organization is a critical first step in understanding your company's risk profile and building an effective remediation strategy.

#2

Restrict Personal Accounts and Enforce SSO

Most GenAI tools now offer corporate accounts, with built-in security measures not found in personal accounts. While specific capabilities between providers, key features typically include private organizational tenant, not shared with other users of the service, not using data for LLM training, and more. However, these benefits depend on users using such business accounts instead of their personal GenAI accounts.

#3

Prompt Users

Security managers often need to strike a delicate balance between security and productivity. This is especially true for GenAI, which many employees use legitimately and effectively. One potent actionable step to limit the risk of GenAI usage it to prompt users with a reminder message when they access GenAI tools. Such a warning message will not restrict users' activities, but it will remind them of organizational policy, risks, and responsible data usage.

#4

Block Sensitive Information Upload

While many organizations allow uploading information to GenAI for legitimate productivity uses, in some cases, restricting the type of data or manner of sharing with GenAI tools may be unavoidable. Therefore, restricting the manner in which data can be inputted into GenAI tools (for example, blocking the pasting of text or blocking file upload) or applying restrictions specifically on data that has been classified as sensitive are effective ways to prevent GenAI data leakage without having to fully block AI tools.

#5

Control GenAI Browser Extensions

Finally, one primary way users often consume GenAI tools is via browser extensions. Such extensions are installed in the browser, automatically tracking and analyzing user activity. While some AI extensions are from reputable publishers and have legitimate uses, for many such extensions, users often don't know who is really standing behind them and what access they have. This is why restricting GenAI browser extensions is crucial in preventing the leakage of sensitive organizational data.

How LayerX Helps Prevent AI Data Leakage

LayerX is an all-in-one, agentless security platform that protects organizations against GenAI data leakage, detects and enforces controls over ‘shadow’ AI apps, and enforces access controls over GenAI usage, with no impact on the user experience.

LayerX natively integrates with any browser, turning it into the most secure and manageable workspace, with no impact on the user experience.

Key LayerX capabilities that help prevent GenAI data leakage:

- Full discovery and visibility into which GenAI websites and SaaS applications are being used in the organizations
- Which users are using each GenAI tool, and whether they are logging in to it using corporate or personal accounts
- Track what data is uploaded to GenAI tools and how it is inputted (e.g., text input, copy/paste, file upload, etc.).
- Easy-to-use policy wizard that enables security administrators to create finely tuned policies to control GenAI usage (for example, preventing developers from pasting code into GenAI prompts)
- Robust enforcement capabilities that don’t just fully allow or completely block GenAI tools, but offer a range of enforcement capabilities, such as:
  - Monitor only
  - Warn user
  - Prevent with an option for the user to bypass by submitting justification
  - Prevent with no option to overrule
- Automatically classify AI browser extensions, assign
- And more!

These capabilities allow security managers to map GenAI usage in the organization, educate users on security risks, enforce usage only of corporate accounts in GenAI tools, prevent uploading sensitive data to GenAI applications, and block risky AI extensions.

To see how we can help you prevent GenAI data leakage without sacrificing productivity, go to http://www.layerxsecurity.com and book a demo today!