LayerX

Browser Security Report 2025

Real-world data and analysis on the hidden risks and security blind spots found in browsers, and how they expose enterprises to AI, identity, and data threats

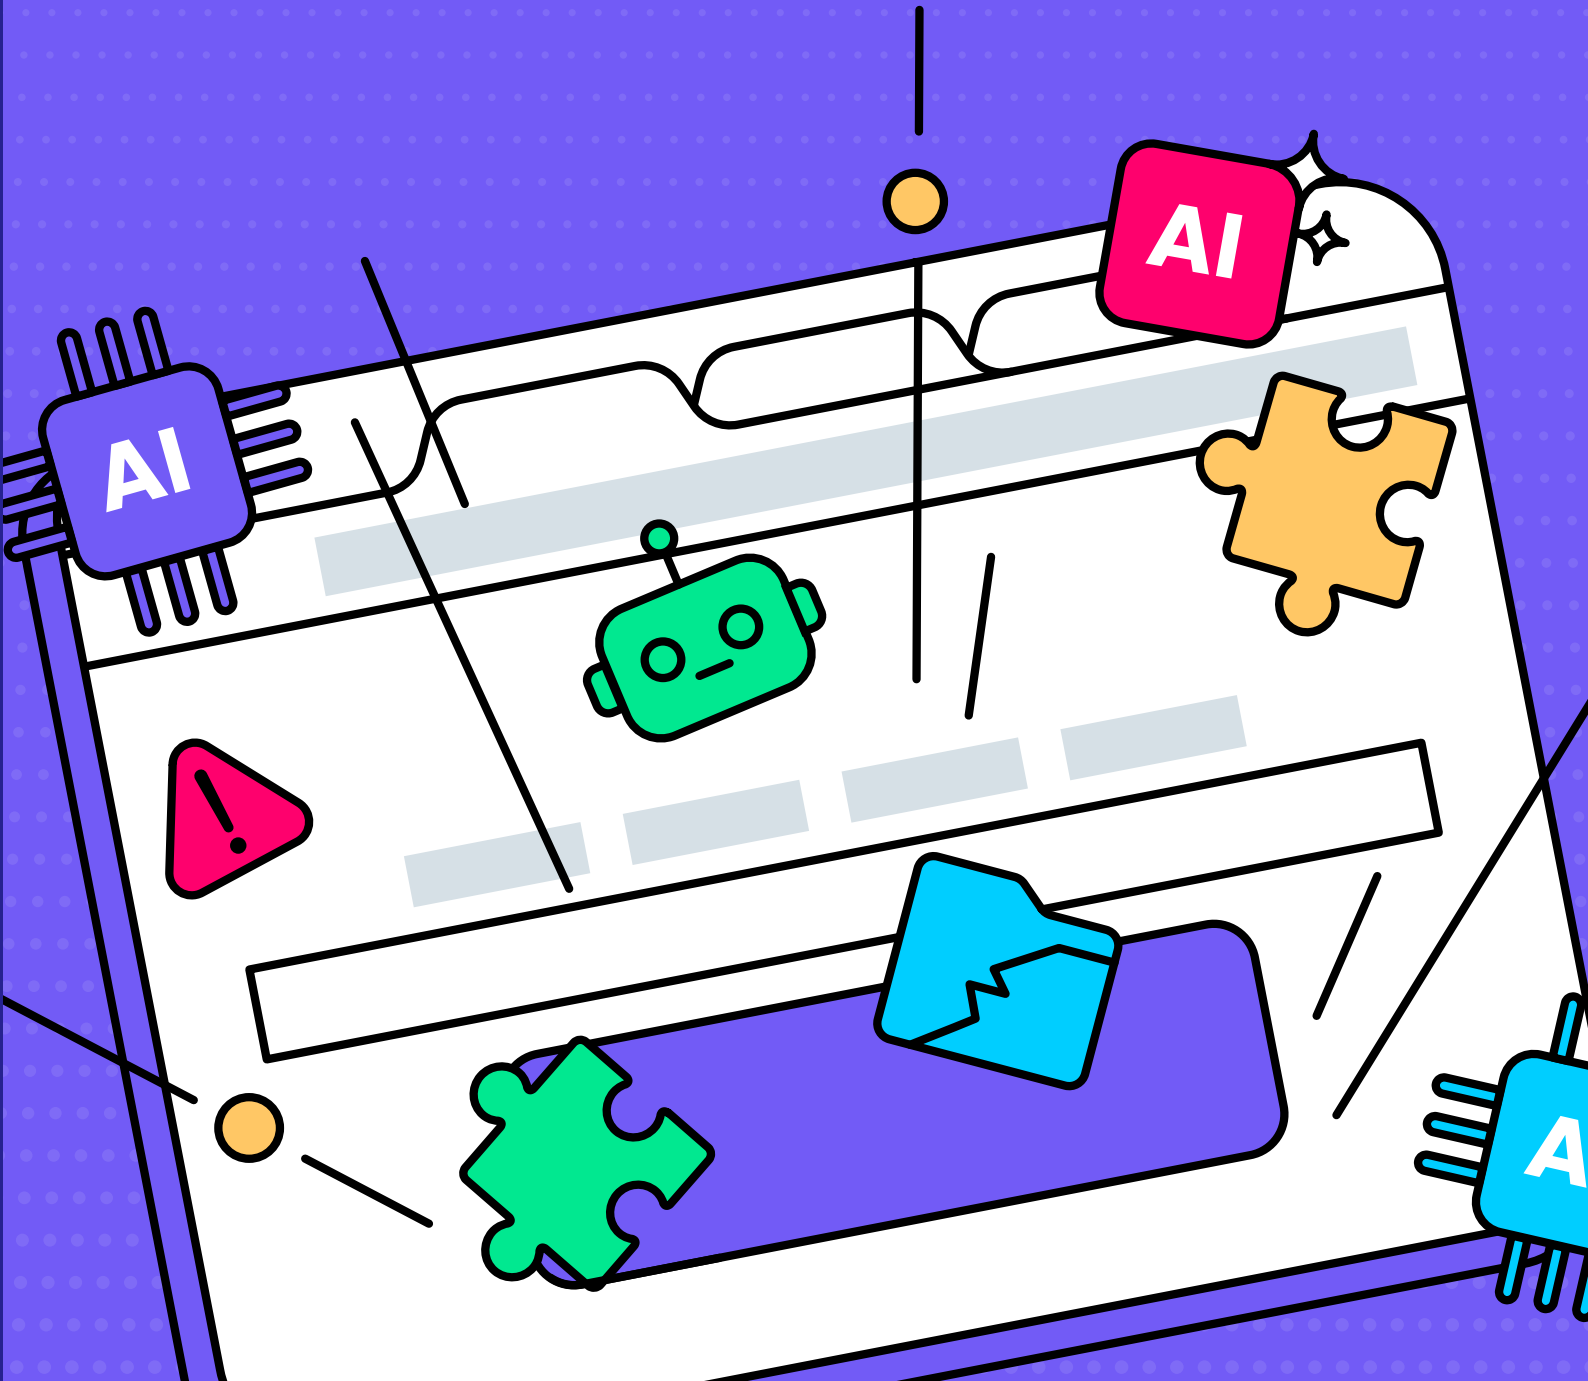

Introduction

Most enterprise work now happens in the browser. It's where employees access SaaS apps, use GenAI tools, authenticate identities, and handle sensitive data. It's where unstructured, last-mile interactions occur. The browser has quietly become the enterprise's most critical and most vulnerable endpoint. Yet despite being at the center of every workflow, the browser remains largely outside the visibility of traditional security stacks like DLP, EDR, and SSE. That blind spot fuels data leakage and credential theft across AI, SaaS and web channels, shadow AI/SaaS, malicious extensions, zero-hour web attacks, and identity risk.

The **Browser Security Report 2025** explores how the enterprise browser has evolved from a simple access point into the most critical layer of the modern security stack. Drawing on telemetry from millions of enterprise browser sessions, the report maps the new risk surface created by the convergence of AI tools, browser extensions, SaaS apps, and unmanaged identities, and the invisible data flows that connect them. It examines the major shifts redefining enterprise security in 2025, from the explosive growth of enterprise AI usage and the rise of browser-based Shadow IT, to the identity and data governance challenges emerging inside everyday workflows. It includes a wealth of statistics, a zoom-in analysis of notable browser-related attacks and trends.

For CISOs and security leaders, this report offers more than just data, it provides a roadmap. It pinpoints where enterprise controls are failing and where modern breaches truly begin. It shows how the security perimeter has shifted from devices and networks to the browser itself. It also outlines how organizations can secure this new control plane for data protection, productivity, and compliance.

Executive Summary

#1

AI Is the Fastest-Growing and Least-Governed Data Channel

Nearly half of all enterprise employees now use GenAI tools, but 90% of sessions happen outside IT oversight. 77% of users paste data into prompts, with 82% of that activity occuring from personal accounts, and 40% of uploaded files containing PII/PCI data. With GenAI now responsible for 32% of all corporate-to-personal data movement, it has become the #1 exfiltration channel in the enterprise browser.

#2

Browser Extensions Have Become the Enterprise's Largest Unmanaged Supply Chain Risk

Over 50% of the extensions installed by enterprise users have 'high' or 'critical' permissions. Even more alarming, 54% of publishers use free Gmail accounts, and 26% of extensions are sideloaded, creating a massive blind spot of unvetted code operating inside enterprise browsers. This isn't just a visibility issue, it's a software supply-chain crisis in disguise, where unverified, abandoned, and AI-powered extensions have full access to enterprise data inside the browser and none of the oversight.

#3

Enterprises Are Blind to Most Identity Usage, Turning Critical Business Apps into Shadow Accounts

Over two-thirds of corporate login events are done without SSO. Moreover, over 40% of SaaS applications in organizational networks are accessed via personal accounts. Identity governance often stops at the IdP but the real exposure begins inside the browser session, where credentials, cookies, and access tokens circulate unprotected. This means security and IT teams are blind to usage of these accounts, and have little-to-no visibility and control over their activities, or where they are used.

#4

While Enterprises Secure Uploads, Most Sensitive Data Leaks Through Copy/Paste, making it the Fastest-Growing Invisible Threat

File uploads are no longer the main channel for data loss. 77% of employees paste data into GenAI tools, and 62% of chat pastes contain PII/PCI, and most of this activity comes from unsanctioned accounts. This means that the majority of data that employees move into GenAI tools is happening outside enterprise oversight, turning copy/paste into a massive blind spot for data leakage. Copy-paste has quietly replaced file transfer as the #1 exfiltration vector for enterprise data, bypassing every file-based DLP control in place.

#5

The Browser Has Outgrown Existing Security Stack to Become the Enterprise's Largest Unsecured Endpoint

From shadow AI to sideloaded extensions and unmanaged logins, the browser now touches every identity, every SaaS app, and every piece of enterprise data. Yet it remains outside the visibility of traditional DLP, EDR, and SSE tools. The browser isn't just where work happens, it's where modern breaches begin and securing it is no longer optional.

Chapter 1:
The AI Security Landscape

> AI: The Fastest Enterprise Adoption Curve With The Largest Governance Gap.

45%
Of employees actively use AI tools

11%
Of enterprise SaaS activity comes from AI tools

92%
Of all AI usage is on ChatGPT

67%
Of employees access GenAI tools via personal accounts

77%
Of employees paste data into GenAI Prompts

40%
Of files uploaded to GenAI apps contain PII/PCI

Findings

LayerX data reveals that 45% of enterprise employees actively use AI tools, with ChatGPT commanding 92% of all activity.

Infact, Generative AI already accounts for over 11% of all enterprise application usage, just behind email (20%), online meetings (20%), and office productivity applications (14%). This underscores how quickly GenAI has joined the ranks of foundational business applications in enterprise environments and accounts for a significant portion of the enterprise users' browsing activity.

Traffic Distribution for Entrprise Application

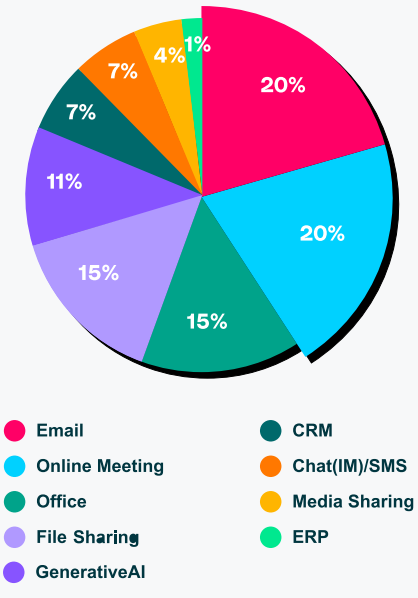

Analyzing AI usage by the number of tools, we see that the vast majority of users (83.5%) use just a single GenAI tool. About 14% use two tools, 2% use three tools, and fewer than 1% of use four tools or more. This distribution falls in line with the finding that ChatGPT accounts for such a high percentage of GenAI tool usage, suggesting that for the majority of users, ChatGPT is GenAI.

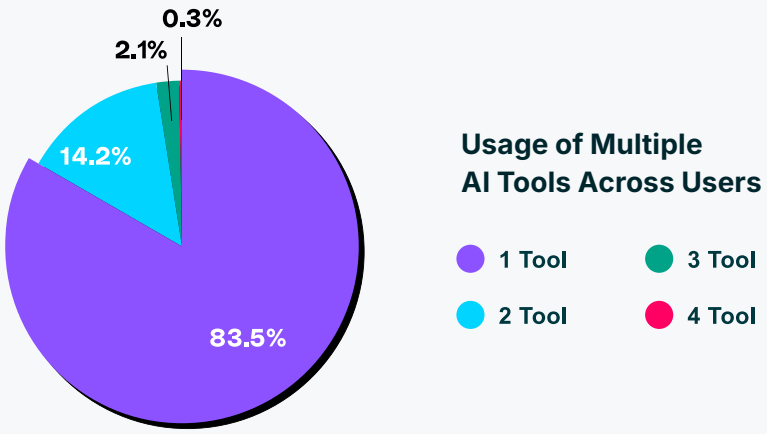

The “top five” AI tools (ChatGPT, Gemini, Claude, Copilot, and Deepseek) account for 86% of AI traffic, but the remaining long tail of smaller AI SaaS tools represents a blind spot of hundreds of unmonitored apps. Few enterprises know which AI tools employees access or whether they comply with data residency or training policies. This growing ecosystem of Shadow AI underscores how enterprise data exposure is no longer limited to mainstream platforms but extends deep into untracked, unsanctioned AI usage across the browser.

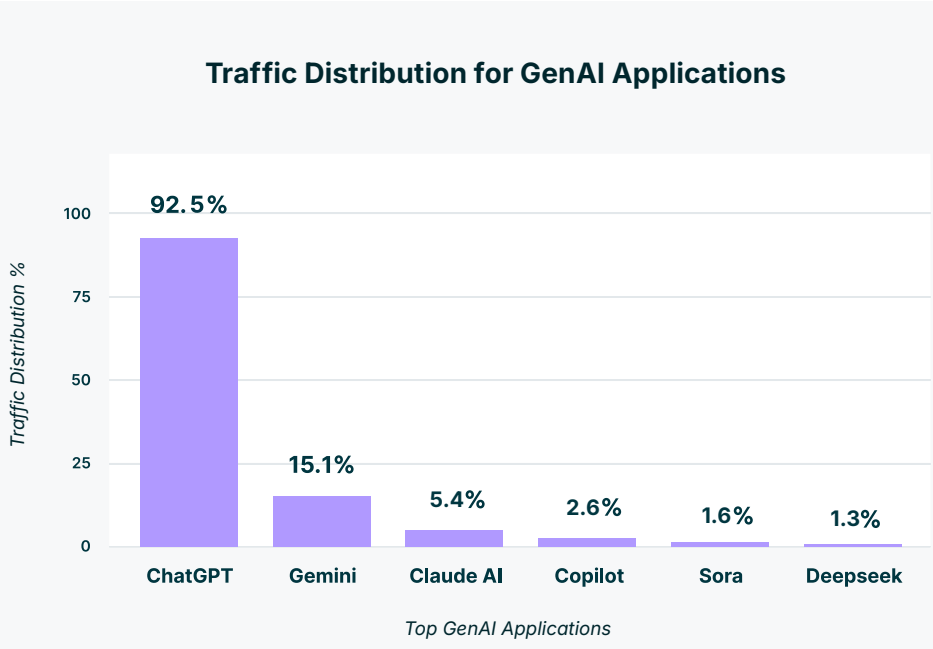

Nearly 90% of AI logins bypass enterprise oversight either through personal accounts (67%) or corporate accounts not backed by SSO (21.3%). Only 11.7% of all AI access truly meets enterprise authentication standards. This means most GenAI sessions take place outside the organization's visibility, leaving IT and security teams blind to what's being shared, typed, or uploaded.

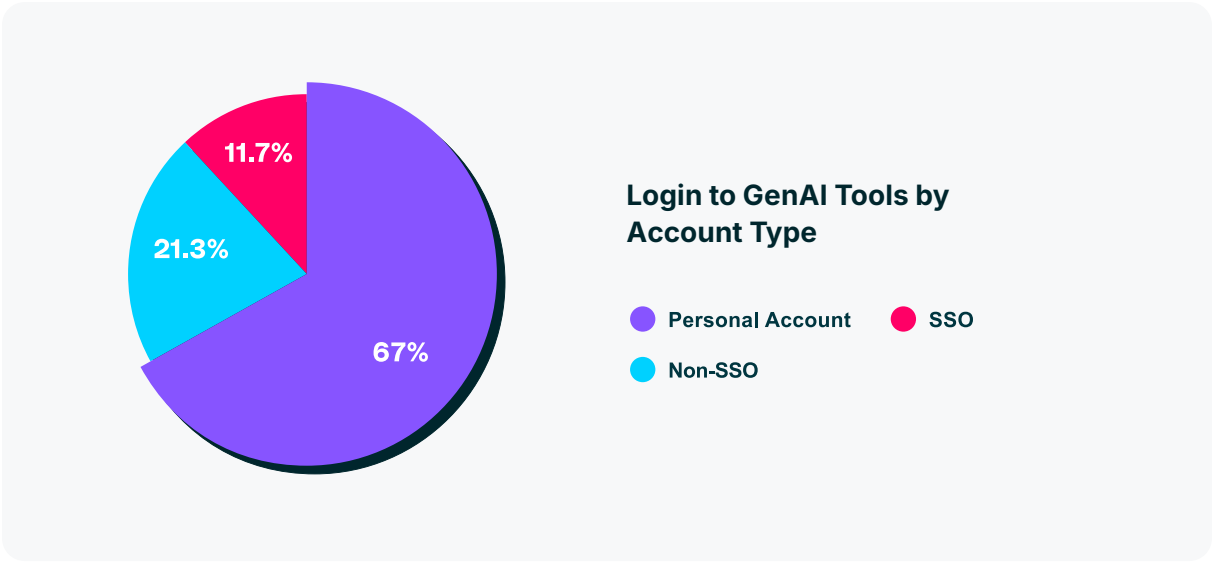

While Email remains the primary file-sharing channel in enterprises, it's not surprising that GenAI apps have now become major upload destinations, with 25% employees uploading files to them. The risk is not just volume, but sensitivity. Over 40% of files uploaded into GenAI tools contain PII or PCI data, turning these platforms into high-risk exfiltration hubs where even a single misstep could result in large-scale breaches and compliance violations.

Beyond files, copy/paste has emerged as the primary data exfiltration channel, bypassing file-based DLP entirely. GenAI tools dominate this behavior as 77% of employees paste data into them, with 82% of that activity occurring via unmanaged personal accounts. GenAI accounts for 32% of all corporate to personal data exfiltration, making it the #1 vector for corporate data movement outside sanctioned enterprise environments.

Analysis

GenAI has become the fastest-adopted enterprise technology in history. They have shifted from "innovation experiments" to the heart of enterprise workflows, yet visibility hasn't caught up.

Nearly half of all employees now rely on AI tools, yet 9 in 10 sessions occur outside enterprise oversight. This creates a massive visibility gap where sensitive data, including PII and PCI, is routinely shared through unmanaged accounts and prompts.

Nearly all enterprise AI activity now happens through the browser, not through installed apps or local clients. Whether it's ChatGPT, Gemini, or Copilot, employees access these tools in-browser, making the browser the new AI endpoint. This shift concentrates both innovation and risk in a single workspace where traditional security controls like DLP and CASB have limited visibility.

The fact that the majority of GenAI usage occurs outside managed environments exposes a fundamental gap between adoption and control. Employees are not malicious; they're simply using AI as a productivity shortcut. But every copy/paste into ChatGPT, every upload to a personal Gemini account, represents potential exposure of sensitive data to public LLMs.

The overwhelming reliance on ChatGPT underscores how quickly a single AI tool has become embedded in enterprise workflows, rivalling long-established SaaS categories. For CISOs, this concentration of use creates both an opportunity to focus governance efforts and a risk that sensitive data funnels into a single, high-volume platform outside of corporate oversight.

Traditional controls built for files and sanctioned apps can't protect against this new form of browser-based data movement. Copy/paste and prompt inputs have replaced attachments as the dominant exfiltration paths, turning the browser into the true frontline of data protection.

Finally, the rapid growth of 'Shadow AI' means exposure isn't limited to major tools, but it's multiplying across hundreds of unsanctioned AI apps that operate beyond IT's line of sight.

For CISOs, this chapter highlights a paradox: the AI adoption curve outpaces governance, making the browser both the launchpad for productivity and the leak point for sensitive data. GenAI governance must now evolve from reactive policy to proactive, monitoring all AI interactions, blocking uploads of classified data, and enforcing SSO-backed access across every browser session. AI security is no longer a niche category; it's an enterprise baseline.

Spotlight:
AI Browsers Are Today's Biggest Enterprise Blindspot: Invisible, Intelligent, and Ungoverned

What Happened

A new wave of browsers, such as Perplexity Browser, Arc Search, Brave AI, OpenAI's Atlas, Dia, Copilot-mode Edge, are blurring the line between browsing, searching, and prompting. These AI browsers embed large language models directly into the browsing experience, continuously reading, summarizing, and reasoning over web pages, tabs, and even local context to "assist" the user. They don't treat AI as an add-on. They make it the operating system.

But what makes them powerful for users also makes them dangerous for enterprises: they act like an always-on AI co-pilot with access to the same session data, cookies, credentials, SaaS tabs, and enterprise content the user can see, but with no enterprise controls or visibility around what that AI collects, stores, or sends back to its cloud model.

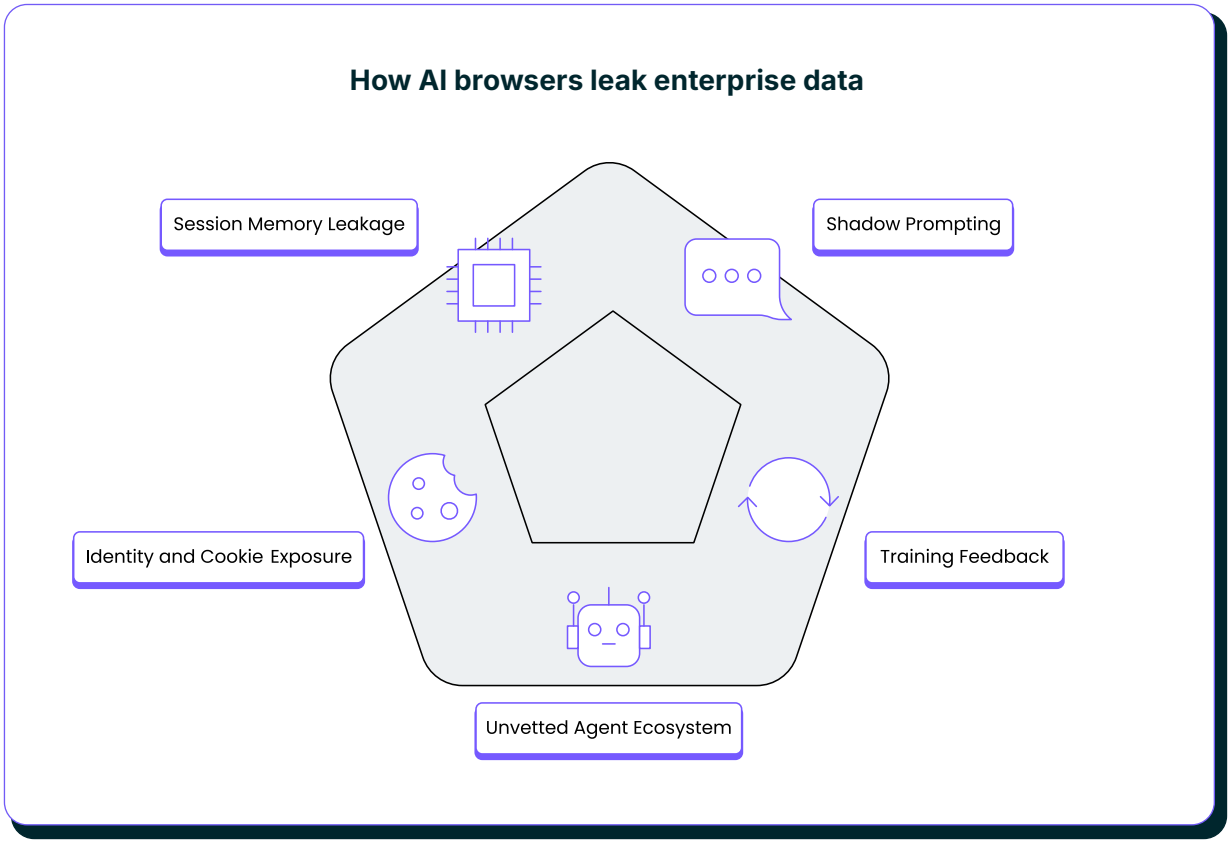

How They Introduce New Data Leakage Vectors

#1

Session Memory Leakage

AI browsers capture session context such as active tab content, search history, copy-pasted data to personalize results.

→ Sensitive information (customer data, financials, code, documents) can be embedded into prompts and sent to the vendor's LLM API.

#2

Shadow Prompting and Auto-Context Injection

Many AI browsers "auto-prompt" the model behind the scenes (e.g., summarize this document, improve this draft).

→ These hidden prompts transmit page content and text selections outside enterprise visibility, a file-less exfiltration path.

#3

Model Training Feedback Loops

Unless explicitly disabled, user activity and content can feed vendor models to improve quality.

→ Corporate session data can indirectly become training material for external AI systems.

#4

Identity & Cookie Exposure

AI browsers use unified login sessions and shared cookies to fetch user-contextual answers.

→ This increases the attack surface for session hijacking and identity replay, especially when AI sidebars interact with authenticated SaaS apps.

#5

Unvetted Plugin & AI Agent Ecosystems

Some AI browsers allow third-party AI "agents" or extensions to automate web actions.

→ These agents can scrape, post, or extract enterprise data with no audit trail or policy enforcement.

Why There's No Oversight

Vendors like Island and Palo Alto's SEB offer a tightly controlled browsing experience. But their model depends on replacing the user's default browser entirely. If a user installs Comet or Dia, these security platforms don't come along for the ride.

In other words: **If you're not using their browser, you're not protected.**

As a result, AI browsers become invisible conduits of data loss and a massive blindspot where sensitive SaaS, identity, and GenAI activity converge. That's simply not realistic in a world where users want to explore the next wave of AI tools.

Chapter 2:
Extension Security and the Hidden Risk Surface

> The Trojan Horse of the Browser: Extensions Slip Past Every Security Layer and Turn into the Browser's Invisible Blindspot

99%
Of enterprise users have at least one extension installed in their browsers

95%
Of Chrome extensions have fewer than 10,000 installs

53%
Of employees installed extensions with 'high' or 'critical' permissions

26%
Of extensions installed in enterprise users' browsers are sideloaded

54%
Of extensions are identified by a free Gmail account

6%
Of GenAI browser extensions are classified as 'malicious'

Findings

A new wave of browsers, such as Perplexity Browser, Arc Search, Brave AI, OpenAI's Atlas, Dia, Copilot-mode Edge, are blurring the line between browsing, searching, and prompting. These AI browsers embed large language models directly into the browsing experience, continuously reading, summarizing, and reasoning over web pages, tabs, and even local context to "assist" the user. They don't treat AI as an add-on. They make it the operating system.

Enterprise Users by Number of Extensions Installed

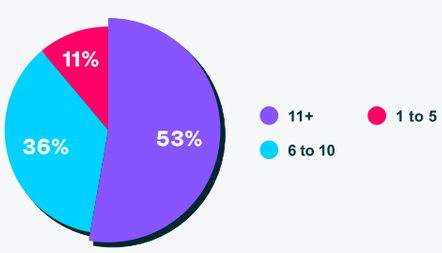

Permission Risk of Extensions Installed by Enterprise Users

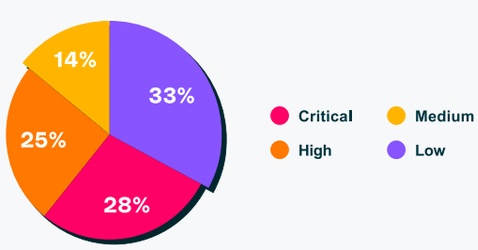

AI-enabled browser extensions are a 'side door' for AI usage in the organization that can often bypass network-layer AI access controls. More than 20% of enterprise users have an AI extension installed, 58% of AI extensions have 'high' or 'critical' level permissions, and 5.6% of them are malicious, making AI extensions an outsized risk should they become compromised. This is particularly concerning as AI extensions typically access sensitive data such as identities, cookies, scripting permissions, and control over browser tabs at twice the average rates of other extensions, blurring the line between legitimate productivity and covert data collection.

Extensions with Access to User Identity

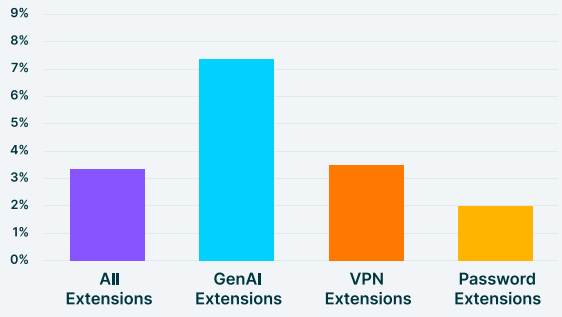

Provides access to user account information when interacting with Google services, including OAuth authentification and access to user profile data.

Extensions with Access to Cookies

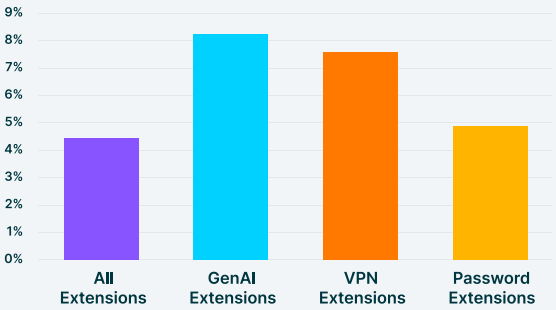

Allows extensions to read, modify and delete cookies. A malicious extension could steal or delete session cookies, effectively hijacking or terminating the user's session and/or create fake cookies to impersonate legitimate users.

While the majority of extensions come from official stores, 17% of extensions originate from non-official stores, and 26% of extensions are sideloaded by external applications, making extensions not just a browser risk but a malware risk, as well.

When discussing browser extension security, extension permissions rightly take up a substantial part of the conversation in an effort to understand what data the extension can access. The second part of the question is how well can I trust it. However, while analyzing permission is pretty straightforward, establishing the trustworthiness of extensions is virtually impossible: 89% of extensions have fewer than 1,000 users, 54% of extension publishers are identified solely by a free webmail account, meaning there is little-to-no information to go by to establish credibility.

Moreover, nearly half (51%) of extensions haven't received an update in more than 12 months. Not only does this open extensions up to software vulnerabilities and supply-chain risks, but it also raises the risk of abandoned extensions that no one is maintaining. Almost 25% of extensions haven't received an update in a year, and are published by a Gmail account, raising the possibility that these are 'hobbyist' extensions that have been abandoned.

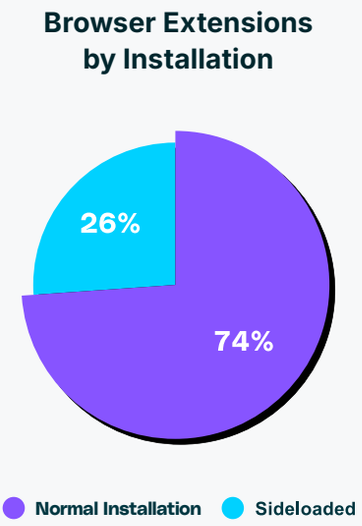

Analysis

Browser extensions have become the most pervasive yet least governed component of enterprise browsing. With 99% of users running at least one extension, they represent an expansive, user-driven ecosystem operating outside IT visibility. While seemingly benign, extensions often hold permissions equivalent to malware. More than half of these extensions hold high or critical permissions, granting access to cookies, session tokens, and in some cases, full browser activity.

AI-enabled extensions take this risk further. Often installed to boost productivity, they request elevated privileges to read page content, capture inputs, and interact with GenAI tools, effectively bypassing network-level AI access controls. The convergence of GenAI extensions and weak identity controls makes them the browser's "backdoor threat."

Many extensions haven't received updates in over a year, are published using free Gmail accounts, or are sideloaded into the enterprise users' systems. This reveals a serious trust gap, making them vulnerable to hijacking and supply-chain compromise, where extensions can act as data exfiltration channels. Together, these trends turn the browser into an unmanaged software supply chain embedded within every endpoint.

For CISOs, the implication is clear: extension security is no longer about permissions. It's about trust, visibility, and control. Traditional endpoint and network tools cannot detect or block these add-ons effectively. Therefore, the browser must now be treated as the main point of risk for extension security that demands continuous monitoring, permission auditing and policy enforcement.

Spotlight:
The Cyberhaven Extension Attack – When Security Tools Become the Threat

What Happened

In December 2024, attackers compromised Cyberhaven's official Chrome Web Store developer account through a consent phishing campaign. The attacker sent a fake Google policy email that tricked a developer into granting OAuth permissions to a malicious app with no passwords or MFA required. With control of the developer account, the attacker uploaded a tampered version of Cyberhaven's Chrome extension (v24.10.4), which was automatically pushed to all users through Chrome's auto-update mechanism.

The malicious version injected new scripts that monitored visits to sites like Facebook, and exfiltrated session tokens, cookies, and account data, effectively hijacking user sessions. Cyberhaven detected and rolled back the update within an hour, but by then, over 400,000 users were already affected. The Cyberhaven incident has therefore exposed a fundamental weakness in how browsers handle trust.

Why It Happened

The attack exploited the browser's blind trust model:

Consent phishing > MFA
OAuth consent flows can grant attacker access without triggering MFA or password prompts.

Auto-update abuse
Once a developer account is compromised, malicious updates silently propagate to all users.

Over-privileged extensions
Browser extensions inherently have deep access to session data, cookies, and active tabs, making them a prime target for data theft.

It wasn't just about Cyberhaven; it was about the browser ecosystem's fragile supply chain.

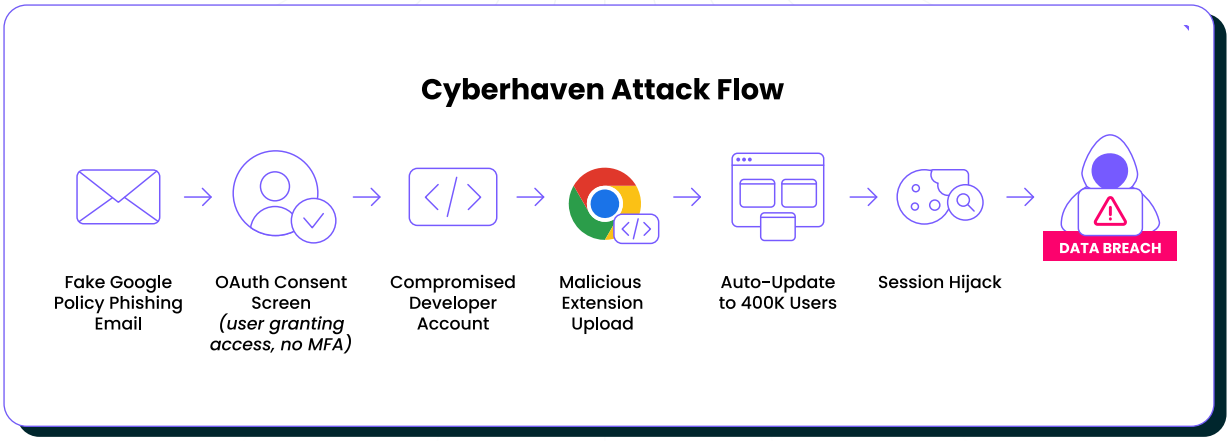

The Implications

The breach exposed a bigger truth: **the browser is now the most powerful and least protected enterprise endpoint.**

- A single compromised extension can instantly turn a trusted security tool into a session hijacking implant.
- Attacks no longer need malware; they hijack trust through OAuth, updates, and permissions.
- Traditional defenses (EDR, SSE, DLP) don't see or control this layer because it happens inside the browser.

Key Takeaway:

The Cyberhaven attack wasn't just a breach. It was a wake-up call proving that in 2025, the battle for enterprise security is happening inside the browser.

Chapter 3:
Identity Security in the Browser

> Identity Security Doesn't End at the IdP;
> It Starts in the Browser

43%

Of logins to SaaS apps are done using non-corporate accounts

68%

Of logins using corporate accounts are done without SSO

67%

Of employees access GenAI tools via personal accounts

54%

Of corporate accounts use passwords of medium strength or below

26%

Of enterprise users re-use passwords on multiple accounts

8%

Of browser extensions in the enterprise have access to the Identity API

Findings

The enterprise identity has become the modern perimeter, but in practice, it's full of holes. LayerX telemetry shows that over two-thirds of corporate login events are done without SSO, and 43% of SaaS applications in organizational networks are accessed via personal credentials. This means security and IT teams are blind to the usage of these accounts, and have little-to-no visibility and control over their activities, or where they are used.

Even though enterprises believe they are securing business-critical apps, account usage tells a different story. Employees are not just using critical apps, they're often accessing them through unmanaged personal accounts. Non-corporate accounts dominate categories like Generative AI (67%), Chat/IM (87%), and Online Meetings (60%).

Even when corporate accounts are used, they are often password-based and bypass SSO entirely, leaving massive blind spots. ERP shows 83% of logins without SSO, CRM shows 71%, and File Sharing shows 47%, making these logins functionally equivalent to personal accounts. It's surprising because these are the very systems housing the most sensitive customer and financial data.

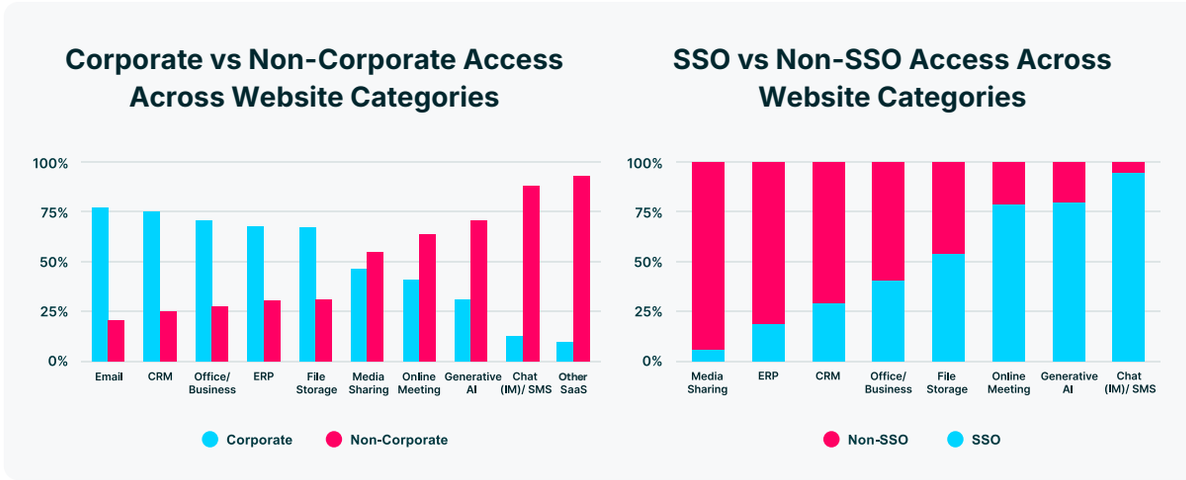

Organizations often assume their Identity Provider (IdP) defines the limits of exposure. In reality, users create shadow identities every day by logging into corporate SaaS apps with personal accounts or non-federated credentials. The result: the enterprise's "identity graph" becomes fragmented, invisible, and unenforceable.

Over 54% of corporate passwords are medium strength or weaker, and 26% of users reuse passwords across multiple accounts. Attackers know this. Any modern password-cracking tools and hardware could easily break them. Credential stuffing remains the most effective breach vector because even strong IAM programs cannot defend against reused personal passwords on unmanaged SaaS accounts.

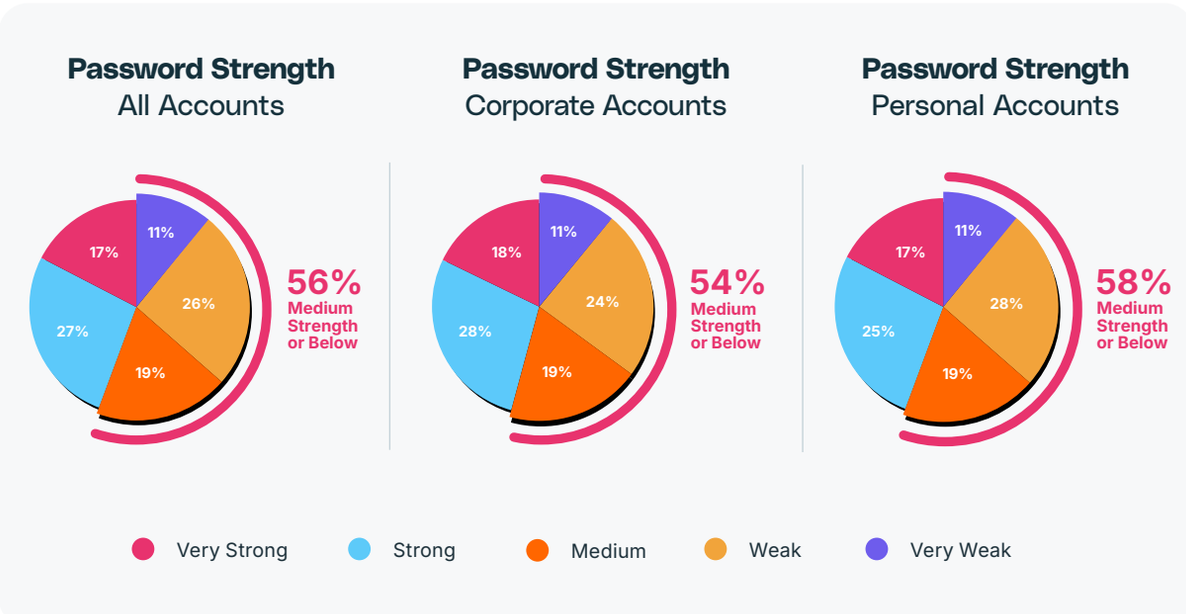

Browser extensions are also a significant threat to users' identity. Nearly 8% of enterprise users installed browser extensions that access their identities, and almost 6% of corporate users have extensions that access their browser cookies. This risk is compounded in corporate environments, where exposure of corporate credentials of one user can lead to a breach that affects the entire organization.

Analysis

Enterprises often believe their IAM stack is secure, yet LayerX data proves that identity exposure happens inside the browser, beyond the reach of the IdP. With 68% of logins occurring without SSO and 43% of SaaS access done through personal accounts, identity governance often stops at the IdP while real risk begins inside the browser. Employees routinely create shadow identities with unmanaged credentials, fragmenting the organization's identity graph and leaving large swaths of activity invisible to IT.

While organizations focus on securing sanctioned applications, usage patterns reveal a different reality, one dominated by unmanaged, non-SSO accounts. Even business-critical systems like CRM and ERP see over 70% of logins without SSO, placing sensitive customer and financial data into environments with minimal oversight. The result is a vast, uncontrolled network of SaaS and AI activity that traditional data protection tools can't reach.

Weak authentication compounds this gap. More than half of corporate passwords are rated medium strength or below, and one in four users reuses passwords across multiple accounts — habits that make credential stuffing and session hijacking alarmingly effective. The browser further amplifies this exposure as 8% of installed extensions can access users' identities and cookies, giving attackers a direct route to harvest credentials and impersonate users within active SaaS sessions.

Beyond logins, user behavior such as copy/paste and autofill drives silent identity leakage. Employees paste data into SaaS and AI tools dozens of times daily, with unmanaged accounts and data containing PII or PCI. These micro-actions create a steady stream of untracked identity data leaving the enterprise, bypassing all perimeter and authentication controls.

For CISOs, the message is clear: identity protection can no longer end at SSO. The browser has become the true identity layer where authentication, access, and data converge. Identity protection needs to move from centralized enforcement to browser-native, real-time monitoring that secures sessions, not just logins. Securing it requires continuous visibility into every session, credential, and extension that touches corporate accounts.

Spotlight:
Scattered Spider - The Attack That Redefined Identity Exploits

What Happened

"Scattered Spider" is a financially motivated, sophisticated threat group known for highly targeted identity and access compromise campaigns. The group gained global attention after successfully breaching multiple large enterprises, including critical infrastructure, telcos, and casinos.

Their hallmark tactic: manipulating the human layer through social engineering to hijack identity access, then moving laterally through browser sessions and SaaS environments.

How the Attack Happened

The modern attack path starts with people and ends inside the browser.

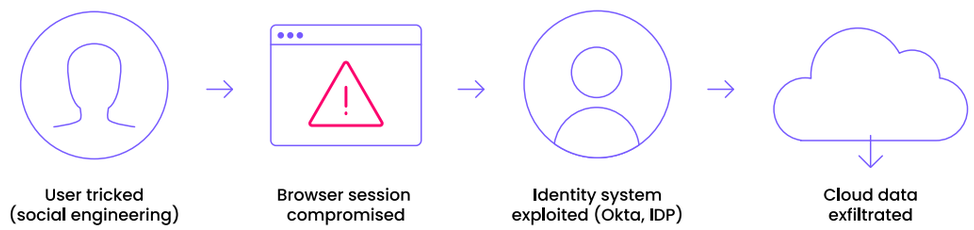

Initial Access via Social Engineering
Attackers impersonated IT helpdesk staff and tricked employees (including contractors) into sharing credentials or resetting MFA.

Session Hijacking
Once inside, they bypassed MFA by stealing valid browser session tokens. This allowed them to log in without needing credentials or MFA re-prompts.

Privilege Escalation & Lateral Movement
Using legitimate accounts and browser sessions, they escalated privileges across SaaS apps, VPNs, and IdPs like Okta.

Data Exfiltration & Disruption
They targeted sensitive data in cloud apps and identity systems and in some cases deployed ransomware to amplify impact.

They exploited MFA fatigue and session token theft to bypass traditional security layers. This wasn't a malware-led breach, it was a browser identity takeover.

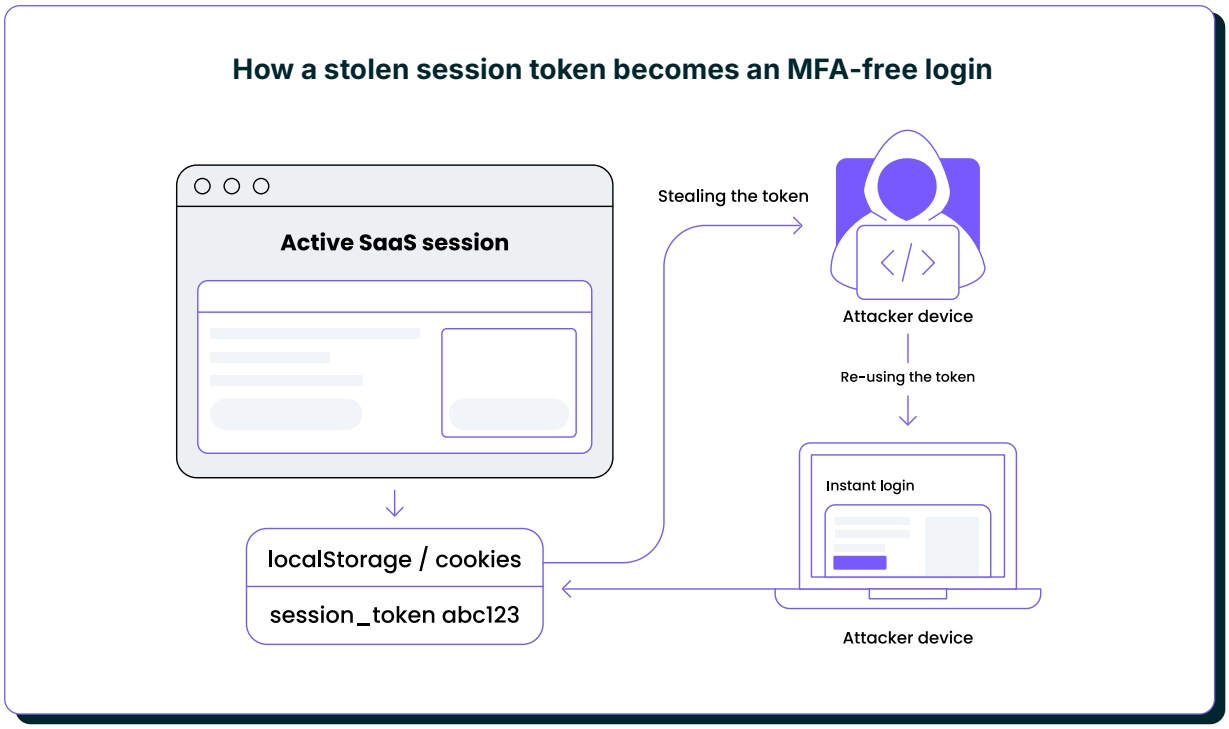

Implications for Enterprises

- **SSO is Not a Silver Bullet:** Even with MFA, attackers can sidestep controls by hijacking live sessions.
- **Browser Sessions are the new "Golden Ticket":** Once hijacked, they give attackers instant, MFA-free access to corporate apps and data.
- **Rapid Lateral Movement:** With legitimate tokens, attackers move invisibly through trusted apps.
- **Traditional Security Gaps:** Network and endpoint tools don't detect in-browser threats like token exfiltration, malicious extensions, or identity drift.

Key Takeaway:

Scattered Spider redefined what "identity exploit" means. It's not about stolen passwords, it's about stolen sessions. The enterprises that win this new battle will be those that see, control, and protect what happens inside the browser.

Chapter 4:
Enterprise SaaS Data Exposure in the Browser

> Your Most Trusted Apps Are Your Least Monitored Data Exfiltration Channels.

38%
Of employees upload files to file storage/sharing platforms

41%
Of files uploaded to file storage/sharing platforms contain PII/PCI

40%
Of files uploaded to GenAI apps contain PII/PCI

62%
Of enterprise users paste data containing PII/PCI into Chat/ IM apps

32%
Of data pasted from corporate to non-corporate accounts is in GenAI tool

4%
Average non -corporate pastes/ day per user containing PII/PCI into SaaS, Office and AI tools

Findings

Workflows remain anchored in communication and collaboration, with Email and Online Meetings dominating enterprise usage, with 80% employees using them. This is followed by File Sharing and Business Applications that remain central to daily workflows as the next most common categories, with 58% of employees relying on them.

File uploads are central to enterprise workflows. It's not surprising that Email remains the primary file-sharing channel, with 64% of employees uploading files to it. However, employees are also moving vast amounts of data not just through sanctioned storage or email but also SaaS, AI, and collaboration tools.

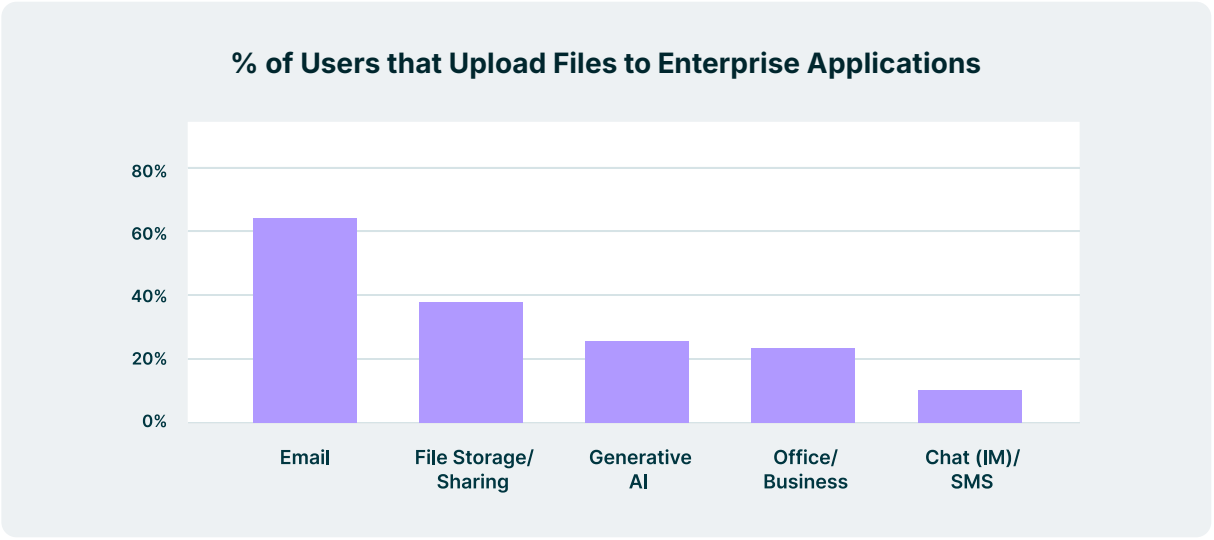

With 38% employees uploading files to File Storage/Sharing tools and 25% to GenAI apps, they have now become major upload destinations. The risk is not just volume, but sensitivity. Over 40% of files uploaded into GenAI tools and 41% uploaded into File Storage platforms contain PII or PCI data. This means that nearly half of the data flowing into these platforms is highly sensitive, turning them into risky exfiltration channels where even a single misstep could result in large-scale breaches and compliance violations.

Copy/paste represents an invisible data leakage stream that existing DLP solutions cannot monitor. GenAI tools dominate this behavior, as 77% of employees paste data into them. File Storage, which accounts for 46% is the second largest paste channel and is followed by Chat/IM and CRM, which are at about 15%. While lower overall in volume, pastes into business-critical apps carry outsized risks because of the nature of the data involved.

However, sensitive data exposure is most severe in Chat/IM, where 62% of pastes contain PII/PCI and 87% of the data is pasted from unmanaged, non-corporate accounts. This makes Instant Messaging apps one of the biggest blind spots for sensitive data leaks. Office apps (20%) and File Storage (17%) come next, with sensitive data being pasted into them frequently.

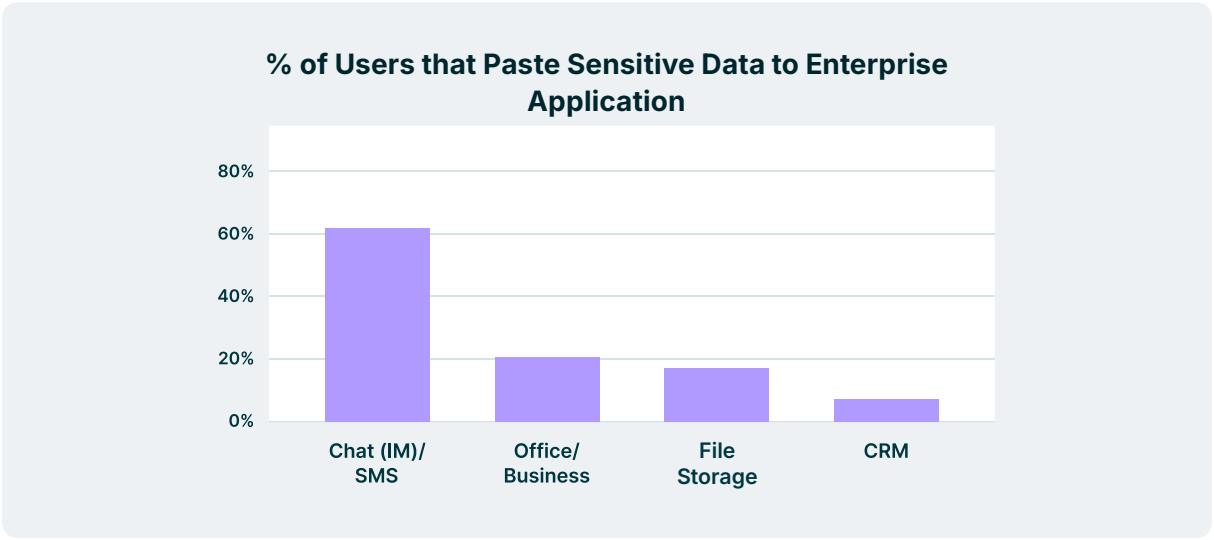

Moreover, GenAI accounts for 32% of all corporate to personal data exfiltration, making it the #1 vector for corporate data movement outside sanctioned environments.

On average, employees make 46 pastes per day. While corporate accounts carry a higher volume of 42/day, non-corp accounts carry a higher risk, averaging 15/day, of which 4 contain sensitive PII/PCI data. This means that even though personal accounts see fewer pastes than corporate ones, they carry a far higher concentration of sensitive data, making them a disproportionately risky channel for invisible data exfiltration.

Analysis

The modern enterprise runs on SaaS, but visibility into how employees actually move data across these apps has all but disappeared. While organizations focus on securing sanctioned applications, usage patterns reveal a different reality, one dominated by unmanaged, non-SSO accounts. Even business-critical systems like CRM and ERP see over 70% of logins without SSO, placing sensitive customer and financial data into environments with minimal oversight. The result is a vast, uncontrolled network of SaaS and AI activity that traditional data protection tools can't reach.

Data movement within this ecosystem is no longer confined to files or attachments. Employees now share sensitive data through uploads, copy/paste actions, and prompt inputs that bypass file-based DLP altogether. Nearly half of all files uploaded to GenAI and File Storage platforms contain PII or PCI data, and 62% of pastes into Chat and IM tools include sensitive information, the majority from unmanaged personal accounts. These invisible streams of data transfer transform routine workflows into continuous, unmonitored exfiltration channels.

This fragmentation of access and control has turned the browser into the true center of enterprise data exposure. It's where employees work, share, and interact, but also where governance ends. The browser sees every upload, paste, and login, yet remains the least protected layer in the enterprise stack.

For CISOs, the path forward is clear: data protection must evolve from app-based coverage to browser-native visibility, where every action, file upload, paste, or prompt is continuously monitored, classified, and controlled in real-time to stop data exfiltration before it happens.

Spotlight:
The Rippling and Deel Scandal – When Messaging Apps Became the New Data Breach Vector

What Happened

In mid-2025, internal messages between executives at Rippling and Deel were leaked, exposing confidential client information, deal data, and internal strategy discussions. The leak originated not from a system hack but from unsecured third-party messaging apps connected to Slack and WhatsApp. These apps, used for sales and recruiting automation, were found storing and transmitting sensitive data to external servers for "training" purposes.

While both companies downplayed the incident as a misconfiguration, forensic analysis showed that the integrations had full read/write access to private message history and attachments. The leak quickly escalated from a "vendor issue" to a broader indictment of how deeply enterprise workflows now depend on unmonitored SaaS and messaging extensions.

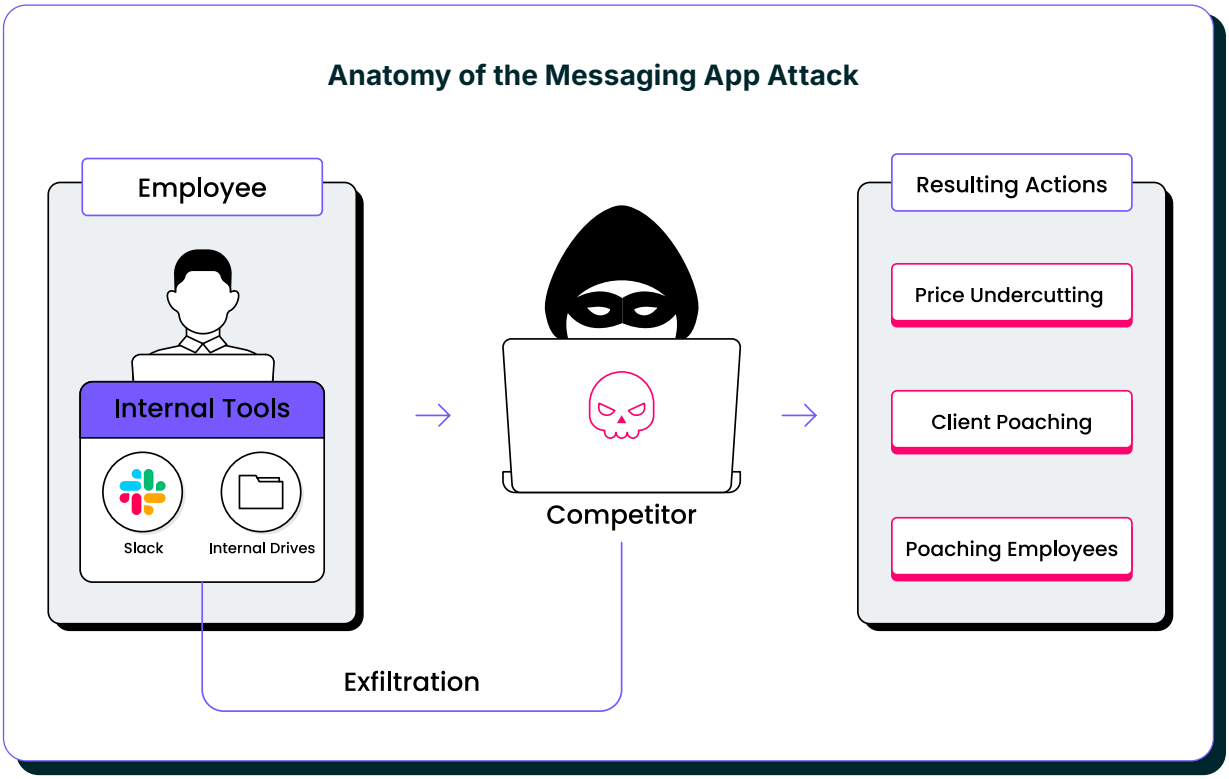

Why It Happened

The incident was a classic case of “shadow automation” with employees enabling AI or chat integrations without security review.

Unvetted app permissions:

Messaging extensions were granted excessive OAuth scopes (read/write/chat: full).

No data residency controls:

Sensitive conversations were copied to non-compliant cloud regions.

Invisible inside browsers:

These integrations ran in browser sessions and APIs completely bypassing SSE, CASB, and DLP tools.

In short, the breach happened because the browser and its connected ecosystem have become invisible to enterprise security oversight.

The Implications

The Rippling-Deel case shattered the illusion that enterprise collaboration apps are “secure by default.” A single unsanctioned plug-in or chat AI can quietly exfiltrate thousands of internal messages with no malware required.

It demonstrated how the browser + SaaS + extension triad is now the most exposed link in the enterprise security chain. Traditional defenses focused on endpoints or networks miss this entirely because the leakage happens inside the user’s authenticated session.

The Rippling-Deel scandal wasn’t about bad actors; it was about blind trust. In 2025, the browser isn’t just where work happens; it’s where sensitive data leaves the enterprise.

Framework For Securing The Browser Against Modern Cyber Threats

The browser has become the control plane for enterprise work and the primary frontier for data, identity, and AI risks. Traditional endpoint or network tools can't protect what happens inside the browser. Use this framework as a practical roadmap to secure your organization's browser environment without disrupting productivity.

A Practical Checklist for CISOs and Security Leaders

Gain Visibility into Browser Activity

- ☐ Inventory all browsers in use: managed, unmanaged, and AI-enabled
- ☐ Map user activity across SaaS, GenAI, and web tools
- ☐ Monitor real-time data movements: copy/paste, uploads, prompts
- ☐ Detect Shadow IT and unapproved browser usage

Govern AI and Extension Usage

- ☐ Maintain allow/block lists for AI tools and extensions
- ☐ Detect Shadow AI activity and non-compliant tools
- ☐ Restrict sensitive data sharing with external AI models
- ☐ Enforce AI and extension risk assessments

Strengthen Identity and Session Controls

- ☐ Enforce SSO and MFA for all browser-based logins
- ☐ Block session hijacking and identity replay attempts
- ☐ Monitor personal vs. corporate account crossover
- ☐ Continuously validate active sessions for anomalies

Protect Data at the Point of Interaction

- ☐ Apply in-browser DLP to monitor file uploads and unstructured data inputs like copy/paste
- ☐ Classify and block risky data actions (e.g., PII in AI prompts)
- ☐ Extend data controls to unmanaged browsers
- ☐ Integrate contextual policies that adapt to user behavior

Secure the Browser Supply Chain

- ☐ Continuously audit extensions, plugins, and AI agents
- ☐ Disable unvetted or automatically installed add-ons
- ☐ Track extension updates and developer reputation
- ☐ Control integration permissions and API access

Manage Browser Risk Continuously

- ☐ Enforce browser configuration baselines and policies
- ☐ Monitor drift in settings, cookies, and permissions
- ☐ Use analytics to detect anomalous AI or data activity
- ☐ Automate adaptive risk scoring and policy response

Empower Users Without Restricting Them

- ☐ Educate employees on AI data risks and browser hygiene
- ☐ Provide a list of safe, sanctioned AI and SaaS tools
- ☐ Align controls with productivity, not against it
- ☐ Adopt a "security without disruption" model

Securing the browser isn't about locking it down, it's about making it self-defending. By combining visibility, identity governance, and data-aware controls directly inside the browser, enterprises can turn their biggest blindspot into their most powerful security layer.

The All-in-One AI & Browser Security Platform

LayerX agentless AI & browser security platform protects enterprises against the most critical AI, SaaS, web and data leakage risks across any browser, application, device and identity, with no impact on user experience

Integrates with All Commercial, AI and Enterprise Browsers

The LayerX Security Platform

Delivered as an Enterprise Browser Extension, LayerX offers the most comprehensive visibility and enforcement capabilities over AI and Browsing risks, including:

AI Usage Security

GenAI DLP
Prevent leakage of sensitive data on AI tools

AI Browsers Protection
Protect AI browsers against attack and exploitation

Shadow AI Discovery
Discover and enforce security guardrails on all AI apps

AI Access Control
Restrict user access to unsanctioned AI tools or accounts

AI Misuse Prevention
Protect against prompt injection, compliance violations, and more

AI Response Validation
Ensure AI response validity and data security

Enterprise Browser Security

Web/SaaS DLP & Insider Threat
Prevent data leakage across all web channels

Browser Extension Management
Detect and block risky browser extensions on any browser

Shadow SaaS & SaaS Security
Discover 'shadow' SaaS and enforce Saas security controls

Safe Browsing
Protect all browsing activity against web exploits

SaaS Identity Protection
Discover and secure corporate and personal SaaS identities

BYOD and Secure Access
Secure SaaS remote access by contractors and BYOD