Quantifying Shadow AI Risk in the Browser

How workforce behavior exposes corporate data — and outpaces the controls built to govern it.

CONTENTS

Inside this report

THE HARD NUMBERS

93.4% work primarily in the browser

70.9% use AI tools daily or more

68.7% pasted sensitive data into AI tools

49.3% would use an unapproved tool if their AI tool of choice was not approved

39.6% knowingly use unapproved AI tools

EXECUTIVE SUMMARY

The AI workplace has arrived. The governance perimeter has not.

The browser is where work happens, where AI is used, and where data moves. It is also where most organizations have the least visibility. This report is built on primary survey research fielded in May 2026 with 227 US knowledge workers who use AI at work, and it focuses on one thing: the choices individual employees make when they open a browser, reach for an AI tool, and decide what to share with it.

The scale of adoption is not the finding — it is the governance gap that shadow AI use has created. As workers reach for newer, faster, better tools to do their jobs, they reach outside company policy to do it. **42.3% already run a governed and an ungoverned AI surface in parallel**, in the same browser session, with the same data.

68.7% pasted internal, sensitive, or regulated data into AI tools in the last three months — financial information, customer data, source code, credentials. For organizations under a regulatory framework, the exposure does not require a breach to be consequential. In many frameworks, the paste itself is the event.

The most consequential finding is not the volume of exposure. It is this: 63.0% of respondents report having a clear AI policy they understand. 48.3% of knowledge workers who know that a policy exists also knowingly breach that policy anyway. **Policy clarity is not the constraint. Enforcement is.**

93.4%

do their work and access business applications primarily through the browser

70.9%

use AI tools daily or more often

68.7%

have pasted or uploaded internal or sensitive data into AI tools

WHAT THIS MEANS

The control point has shifted to the browser — to the moment of the prompt, the paste, the click. Policy without a browser-level control point is just a document waiting to be ignored.

FINDING 01

The browser is the workplace

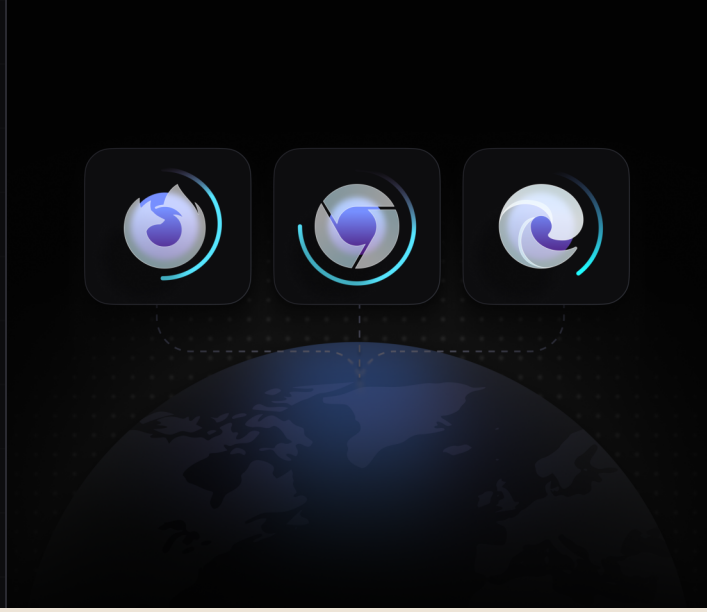

93.4%

do their business application work primarily in the browser. 71.4% work there almost entirely.

That figure is not a trend; it is a settled reality. Every security architecture built around the endpoint, the network perimeter, or the email gateway was designed for a workplace that no longer exists for knowledge workers. The browser session is where company data is accessed and shared, where decisions are made, and where risk gets created.

And these are not clean, controlled environments. Ordinary, habitual choices create a persistent blur between personal and professional identities — between what an organization *thinks* it controls and what it actually does.

THE PERSONAL / WORK BLUR

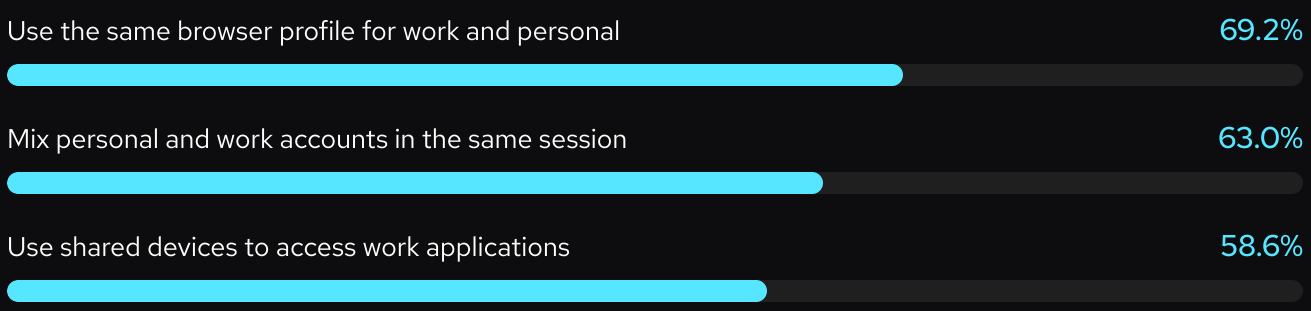

The browser is the workplace. It is also the largest ungoverned surface in the environment.

FINDING 02

AI use is daily, essential and employee-driven

70.9%

use AI tools daily or more. This is not a pilot program.

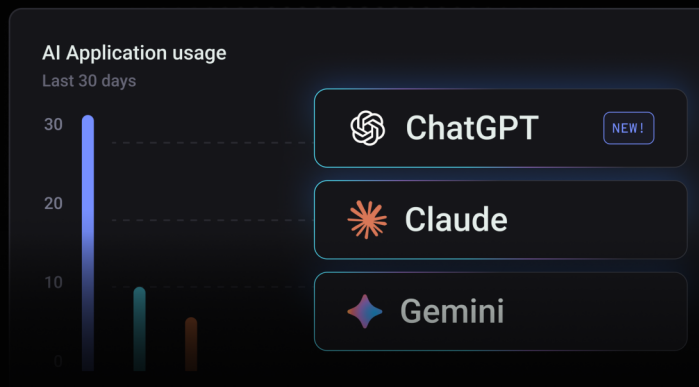

51.1%
use AI tools multiple times a day

63.4%
rate AI as essential or very necessary to their job

32.6%
call it *absolutely* essential

Workers are not using one AI application and stopping there. They build a stack — and the layers do not all sit where IT can see them. *Respondents could select multiple options.*

WHERE AI SHOWS UP IN THE STACK

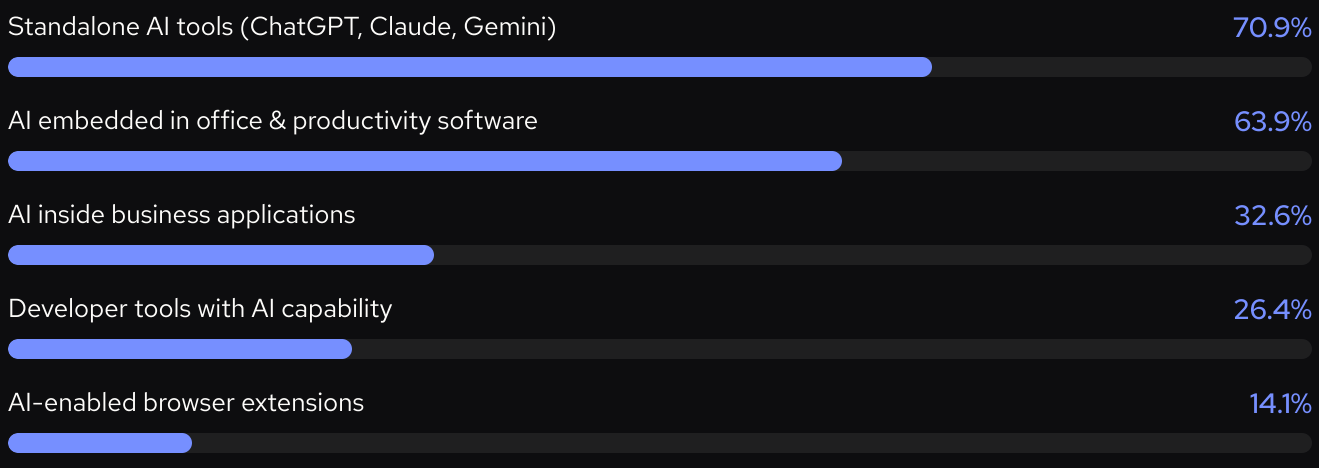

> **42.3% run both a governed office AI and an unsanctioned AI tool in the browser.** IT can see one. The other runs alongside it, invisible — deploying Microsoft Copilot or Workspace AI governs one surface and leaves the other unsupervised.

FINDING 03

More than half work outside sanctioned AI boundaries

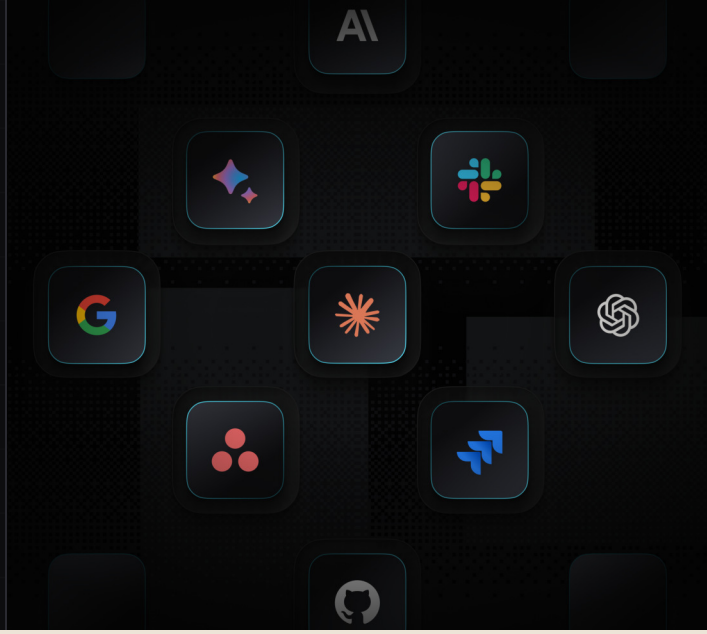

54.6%

use AI tools their organization has not approved.

39.6% said definitively they had used unapproved tools; a further 15.0% said they might have. A worker uncertain whether a tool is approved has not checked — that is itself a governance signal.

Shadow AI has no seniority floor. Among individual contributors, 14% reported knowingly using unapproved tools; among Managers, 35%. The pattern does not reverse as authority and data access increase.[1]

SHADOW AI BY STACK COMPLEXITY

Share reporting use of unapproved tools, by number of AI tool types used

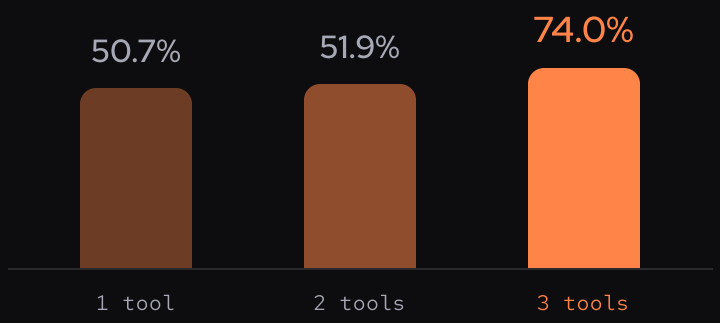

Workers who build a broader AI stack reach further outside policy to do it.

The workaround data explains why policy alone does not hold. Asked what they would do if a wanted tool was not yet approved:

49.3%
would use a different unapproved tool anyway

41.9%
would copy work data into a personal tool "just this once"

Workers are not waiting for the process. They are not breaking the rules — they are outrunning them. This is rational behavior where AI delivers a genuine productivity advantage and approval timelines do not match the pace of work; it may also reflect workers, concerned about layoffs, motivated to build AI skills to stay relevant. **41.4% expected approval to take more than two weeks; 20.3% more than a month.** The gap between what workers need and what IT can provision in time is where shadow AI lives.

[1] Job-level figures are directional. Subgroup sizes are small (IC n=44, Manager n=79, VP n=26, Director n=45) and margins of error range from ±11 to ±20 points. Findings should not be treated as precise prevalence rates within each level.

FINDING 04

Sensitive data is already inside unsanctioned AI tools

68.7%

pasted internal, sensitive, or regulated data into AI tools in the last three months.

Only 21.6% never paste anything sensitive. For everyone else the question is not whether data has moved into AI tools — it is how much, how often, and what kind. Among those who paste, it is a habit, not an exception.

70.5%

of those pasting sensitive data do it habitually — 42.9% almost every time, 27.6% most of the time

The browser is where this happens. Not in email. Not in a file transfer. In a prompt field, at the click of a button — the moment current controls were not built to see.

WHAT IS BEING PASTED

Share of all respondents · multiple selections allowed

General work content	48.5%
Internal business information	30.8%
Financial information	25.6%
Customer data	25.1%
HR information	20.3%
Regulated / sensitive personal data	20.3%
Source code	13.2%
Credentials or API keys[3]	11.0%

For organizations under a regulatory framework, the exposure does not require a breach to be consequential. In many frameworks, the paste itself is the event.

[2] Frequency figures are based on the 156 respondents (68.7% of n=227) who reported pasting at least one sensitive data type. [3] The credentials and API keys figure rests on 25 respondents and should be treated as directional only.

FINDING 05

Policy is not the problem. Enforcement is.

63.0%
have a clear policy they understand

48.3%
of those workers knowingly use unapproved tools anyway

Those two numbers sit in the same dataset, from the same respondents. Policy awareness is not the constraint.

The conventional response to AI risk is to write a policy, communicate it, and train the workforce. The data suggests that these policies already exist in the majority of organizations represented here, but it has not actually reduced the risk. Workers who understand the policy are violating it anyway. The problem is not that employees do not know the rules. It is that the rules have no mechanism of enforcement at the moment they matter.

That moment is the browser session. The prompt field. The paste. The click. No policy document, however clearly written, operates at that level of granularity. No email from the CISO reaches the user at the point of action. By the time the training is complete and the policy is understood, the behavior is already in motion.

34.3% AWARENESS GAP

found policy unclear, unseen or uncertain (27.3%), or said none existed (7.0%). A genuine governance gap.

48.3% ENFORCEMENT GAP

of workers who know a policy exists knowingly breach it anyway. Closing the awareness gap will not solve this.

90% of that subset — workers who knowingly breach a policy they understand — are pasting or uploading internal business information or sensitive data of some kind into AI tools.

Policy without a browser-level control point is just a document waiting to be ignored.

WHAT THIS MEANS FOR SECURITY LEADERS

The question is not whether this is happening. It is whether you can see it.

The five findings describe a workforce that has moved decisively into AI, is operating well outside sanctioned boundaries, and is handling sensitive data in ways most organizations cannot observe. The data suggests it almost certainly is happening. What follows is what prepared leaders and practitioners are doing about it — framed for the three readers who need to act.

FOR THE CISO

Visibility is your posture. Enforcement at the point of action is the gap that training cannot close.

p. 10

FOR THE BOARD

The policy paradox. The question is not “do we have a policy?” but “does it have teeth?”

p. 11

FOR THE PRACTITIONER

The internal business case you have been looking for — and three frames that land.

p. 12

FOR THE CISO

Your posture is only as strong as your visibility into where AI is used.

Legacy security stacks — DLP, email security, endpoint detection — were designed for a threat environment where data moved predictably and the perimeter was definable. The browser has dissolved that perimeter. 93.4% of your workforce does its application work there, and the AI tools they reach for, sanctioned and unsanctioned alike, live there too.

The finding that should concern you most is not the 54.6% using unapproved tools, striking as that is. It is the **48.3% of policy-aware workers who breach it regardless.** That gap cannot be closed with better communication. It requires enforcement at the point of action, in the session, before the data moves.

WHAT PREPARED ORGANIZATIONS PUT IN PLACE

01 **Full visibility across the browser surface**

Into all AI and SaaS usage — not just the tools IT has sanctioned.

02 **Real-time controls at the moment of risk**

That operate in the session rather than logging it and investigating afterwards.

03 **A governance narrative, with evidence**

Audit-ready visibility that you can present to the board.

The finding that matters most at board level is the policy paradox.

Your organization almost certainly has an AI policy — the majority of organizations in this research do. The data shows that policy clarity is not the constraint on safe AI adoption. Enforcement is.

That distinction changes how you frame the risk in the boardroom. The question is not "do we have a policy?" Most do. The question is **"how much of our AI risk is invisible to us?"** For most organizations, the honest answer is unclear.

The reputational and regulatory exposure does not require a breach to be consequential. In many frameworks, the act of pasting sensitive data into an unsanctioned tool is itself a compliance event. The risk is not theoretical and not future-tense. It is current, habitual, and in most organizations invisible.

THE BOARD QUESTION

"How much of our AI risk is invisible to us?"

The organizations that are growing and innovating know they cannot restrict AI adoption. They are governing it — with controls that operate where the workforce actually works, that do not slow users down, and that produce audit-ready visibility.

FOR THE SECURITY PRACTITIONER

The data in this report is the internal business case you have been looking for.

54.6% shadow AI use. 68.7% pasting sensitive data. 74% using three AI tools or more. This is a gap that compliance training alone cannot close. These are not projections — they are self-reported behaviors from a workforce that looks like yours. The challenge is translating that into a conversation leadership will act on. Three frames tend to land.

THE VISIBILITY FRAME

"We do not currently have eyes on what our workforce is doing in the browser. That is where the risk is being created."

THE ENFORCEMENT FRAME

"Our policy exists. Our ability to enforce it at the point of action does not. Those are different problems with different solutions."

THE COMPLIANCE FRAME

"The exposure does not require a breach. In several frameworks we operate under, what is already happening may constitute a reportable event."

Prepared teams are building the case for browser-level visibility and enforcement as the missing layer — not a replacement for what exists, but the control point their current tools were never designed to cover.

WHERE PREPARED ORGANIZATIONS GO NEXT

The harder question is where your own organization sits.

The findings quantify a risk most security leaders already suspect. Four questions separate the organizations that know where they stand from the ones still guessing. They do not have comfortable default answers — but they have answers.

Q1 How much of your workforce's AI use is visible to you today?

Q2 How many of the tools they use fall outside your sanctioned list?

Q3 When sensitive data moves into an AI prompt, do you know?

Q4 When policy is violated at the click, do your controls respond in that moment — or discover it afterwards?

Finding out where you stand is the most important next step a security leader can take right now.

THE NEXT STEP

Find out where your organization sits on the curve.

The Neon Cyber AI Exposure Assessment maps your workforce's real browser-based AI behavior against the findings in this report — shadow AI, sensitive-data movement, and the gap between your policy and what is actually happening.

TAKE THE AI EXPOSURE ASSESSMENT →

Structured. Browser-level. Built from this research.

SEE	UNDERSTAND	ACT
Every AI and SaaS tool in use across the browser surface.	Where sensitive data is moving, and how often.	Close the gap between policy and behavior at the point of action.

How this study was conducted

An exploratory, behavior-focused study commissioned by Neon Cyber to examine how US knowledge workers use AI at work, the security behaviors that accompany it, and the gap between stated policy and observed behavior. It is intended to surface directional patterns, not to size a market.

FIELDING & SAMPLE

Fielded May 11, 2026 via an opt-in US online research panel. Opt-in panels do not independently verify employers or job titles; all role and demographic information is self-reported. References to "VPs," "Directors," or specific industries should be read as respondents who self-identified as such.

FILTERING

From 484 responses we excluded: 19 whose job titles fell outside knowledge work (e.g. transportation, maritime, construction, custodial); 99 who exited after the initial AI question without completing substantive sections; and 21 who do not primarily use a browser for business applications. Excluding non-completers means the sample represents AI users specifically, not the workforce as a whole.

SAMPLE COMPOSITION

Mid-career skew (80% aged 30–60), roughly gender-balanced. 77% Team Lead or above; 11.5% VP or above. Company size spans 201 to 10,000+ employees, concentrated in 201–5,000 (74%). Industry mix: Technology 17.2%, Education 16.7%, Manufacturing 14.1%, Healthcare 13.7%, Financial Services 12.8%. This does not precisely mirror the enterprise security buying population; cuts are directional and unweighted.

n = 227

self-identified US knowledge workers who use AI at work — the final analytical sample

CONFIDENCE & PRECISION

At n=227, approximate margin of error is ±6.5 points at 95% confidence. Subgroup margins widen materially: VP (n=26) ≈ ±19 points; Director (n=45) ≈ ±15 points. Subgroup findings, including higher reported shadow AI among senior respondents, are directional patterns warranting further research, not conclusive results.

SELF-REPORT LIMITATIONS

This survey captures stated behavior, not what organizations can independently observe. Both over- and under-reporting are possible, particularly on shadow AI, data-pasting and disclosure. The 11% who reported pasting credentials or API keys rests on 25 respondents and is directional, not a precise prevalence rate.

WHAT THIS STUDY IS — AND IS NOT

A directional snapshot of self-identified US knowledge workers using AI at work, suitable for understanding patterns and informing strategic discussion. It is **not** a statistically representative sample of the US workforce, the enterprise security buying population, or any specific industry; nor a measure of incident or breach rates. Readers should weigh the findings accordingly.

See every browser session.
Understand the risk.
Act in real time.

ABOUT NEON CYBER

Say “yes” to AI without leaking data.

Your workforce is already using AI—but legacy security tools only see fragments of activity, and “locked-down” browsers only frustrate your team and kill productivity.

Neon Cyber provides the missing “last mile” of security. Our AI-native browser security platform gives you deep visibility into how your workforce uses generative AI—every prompt, input, upload and download. And with context-aware guardrails, you can enforce security where work actually happens: in the browser. Stop AI data leakage and reduce insider risk at the point of intent.

Learn more at neoncyber.com.