

APRIL 2026

Securing AI Usage Starts in the Browser

Gabe Knuth, Principal Analyst

Abstract: Generative AI (GenAI) has become embedded in the workday, and the vast majority of organizations now allow employees to use public GenAI tools.¹ But AI introduces data security challenges that traditional approaches struggle to address. The browser, as both a ubiquitous productivity surface and the primary interface to AI applications, is a natural enforcement point for visibility, data protection, and policy. Recent Omdia research confirms that GenAI application security is the leading use case driving secure browsing adoption. Google Chrome Enterprise provides the discovery, enforcement, and governance capabilities organizations need to secure AI usage without sacrificing the productivity it enables.

Introduction

The browser has become the centerpiece of the end-user computing environment at the same time that GenAI usage has exploded across the enterprise. To better understand how organizations are navigating this intersection, Google partnered with Omdia to survey 400 IT and cybersecurity professionals across North America on the state of browser management and security. The results confirmed that AI-driven data protection is a front-burner issue, and that the browser is where organizations expect to address it.

GenAI has changed the browser security equation

The browser's role in the modern workplace is hard to overstate. Omdia research found that, on average, IT practitioners believe the typical knowledge worker spends more than half (56%) of their workday in a browser. Google Chrome is widely adopted, with 88% of respondents indicating that it's formally supported in their organization. Support for other major browsers is also common across the majority of organizations, with many in use but in an unsupported manner. And though Windows apps persist in rather large numbers, it's clear the browser has become the central productivity surface for the modern workforce.


It is also increasingly where GenAI happens. Nearly all organizations (92%) now allow employees to use public GenAI applications, and much of that usage takes place in the browser. Unlike traditional application access, where data goes in, gets processed, and stays within enterprise boundaries, GenAI interactions are designed around sharing data for analysis. That creates a fundamentally different risk profile—one that sits squarely in the browser.

Organizations are well aware of the exposure. When asked about emerging threats, data leakage via unsanctioned AI applications ranked among the top concerns, with 92% of respondents expressing concern. AI-powered phishing and social engineering (95%) is the only threat rated higher, underscoring how central AI-related risks have become to the security landscape (see Figure 1).

¹ Source: Omdia Research Report, *Browser Management and Security: Emerging Strategies, Requirements, and Success Factors*, May 2026. All Omdia research references and charts in this Showcase have been taken from this report.

Figure 1. Data leakage via unsanctioned AI and AI-powered phishing and social engineering attacks top the list of browser-borne concerns

How concerned is your organization about the following types of emerging threats? (Percent of respondents, N=400)



Source: Omdia

What’s more, these concerns are not hypothetical. Among organizations that experienced browser-related security incidents, data loss or leakage was the second most common attack type (38%), second only to phishing (40%). Data is already leaving organizations through the browser, and GenAI is amplifying the problem.

Legacy web security tools, including secure web gateways (SWG), Cloud Access Security Brokers, and network data loss prevention (DLP), have always addressed data egress to some degree. But GenAI data sharing happens within browser sessions in ways that network-level inspection struggles to see: prompt inputs, copy/paste into form fields, and file uploads to legitimate SaaS domains, all over encrypted connections. These are interactions with sanctioned, known-good web properties, making traditional URL-based filtering largely inadequate for an expanding threat landscape.

“Put simply, organizations want visibility into AI activity, integrations with other platforms and processes, data controls that operate at the browser level, and the ability to set and enforce granular policy.”

- Gabe Knuth

Since the browser is where that activity occurs, it makes sense to focus on browser-level controls that offer visibility that network-centric approaches cannot easily match.


A browser-first approach to AI security

Thus far, organizations have turned to a variety of legacy tools, with SWGs (58%), secure browsing solutions (57%), and SaaS security tools (57%) all clustered at the top and no single approach emerging as the clear primary choice. Most of these were not designed with GenAI data flows in mind, and the overlap suggests organizations are still assembling the right combination of controls.

It is clear that organizations want more, and when asked about the most important use cases for a secure browsing solution, GenAI application security ranks first at 59%, ahead of DLP and data security (51%) and general web security enhancement (42%, see Figure 2).

Figure 2. GenAI application security and DLP outweigh even general web security and unmanaged device access when determining secure browsing solution use cases

What use cases are, or will be, most important for a secure browsing solution at your organization? (Percent of respondents, N=400, three responses accepted)



Source: Omdia

When asked about the top attributes for a secure browsing solution, respondents prioritized controls over GenAI application usage (53%), centralized policy enforcement (52%), and data protection and DLP (52%), with integration with other security tools leading overall at 57%. Put simply, organizations want visibility into AI activity, integrations with other platforms and processes, data controls that operate at the browser level, and the ability to set and enforce granular policy.

AI policy and enforcement vs. blocking

Organizations face a real challenge balancing allowing AI usage based on company policy (including enforcement) with usage blocking, which can lead to shadow AI. Frontier and domain-specific GenAI models are extraordinarily capable, but using them with confidential or privileged information raises legitimate trust

concerns. Organizations that block GenAI outright avoid the risk but also forfeit the value. Knowledge workers are already using these tools to draft, summarize, analyze, and automate routine tasks, sometimes without corporate knowledge or support, and replicating those capabilities with internally sanctioned alternatives that run at a fraction of the scale is a difficult proposition.

What organizations need is a way to say “yes” but with conditions. By enforcing conditional, posture-based policies at the point where users interact with GenAI applications, browser-level security tools can allow usage when the right conditions are met and restrict it when they are not, without relying on network-layer controls that lack the context to make those distinctions.

How Google Chrome Enterprise helps

Google positions Chrome Enterprise as a centralized enforcement point for AI security within the browser, offering granular, context-aware controls structured around three capabilities: discovery, enforcement, and governance. The approach works across ChromeOS, macOS, Windows, Linux, and mobile, with no additional agent or network proxy required. That cross-platform, agentless model addresses a real concern for organizations managing diverse device fleets, especially those supporting unmanaged, bring-your-own-device, or contractor devices where installing additional software is not always feasible.

Discovery and visibility

Before organizations can secure AI usage, they need to see it. Chrome Enterprise provides this visibility through Security Insights, which monitors content transfer activity (uploads, downloads, copy/paste, prints, etc.) to identify potential data exfiltration, including transfers to AI application domains. It surfaces high-risk users and domains based on content transfer patterns. Coupled with comprehensive event logging that records security events such as sensitive data transfers, URL filtering triggers, and content transfers flagged by DLP rules, Chrome Enterprise gives IT the ability to better identify shadow AI usage and data exfiltration.

Enforcement and data protection

Chrome Enterprise Premium extends beyond visibility into active policy enforcement. Organizations can block access to unsanctioned AI tools outright or steer employees to approved alternatives. Granular DLP controls at the browser level govern uploads, downloads, copy/paste, print, and screen capture to and from AI application domains, blocking actions and notifying users when sensitive data is detected. Context-aware access for SaaS and web applications uses 30+ device and partner signals to make access decisions based on real-time posture rather than static allow/block lists.

Governance and conditional access

Rather than treating all GenAI usage as a single policy decision, Chrome Enterprise lets organizations apply different rules for different AI tools, user groups, and risk levels. Elevated security posture requirements for AI access allows organizations to permit a sanctioned GenAI tool only when the device has current OS patches, is on the corporate network, and meets management compliance requirements. If conditions are not met, access is denied or restricted. Outbound encryption at the browser level enforces data protection controls at the point of interaction rather than relying on downstream network inspection.

This approach means organizations do not need to implement the full enforcement stack on day one. They can start with Security Insights to baseline AI usage and data exposure, then layer on additional enforcement and governance controls as policies and use cases mature.

Conclusion

As browser usage increases and GenAI becomes more established in the hands of end users, IT and security teams are faced with enabling productivity without increasing data loss or compliance risks. The traditional mix of legacy, network-centric approaches offers coarse-grained control that struggles to distinguish between productive AI usage and risky data sharing and is holding companies back. Our research makes clear that organizations recognize that GenAI security is the top use case driving secure browsing adoption, and they want solutions that provide visibility, policy enforcement, and data protection at the browser level.

Google Chrome Enterprise addresses this by bringing discovery, enforcement, and governance capabilities directly into the browser organizations already use. The graduated approach, from free visibility through Security Insights to full enforcement and governance with Chrome Enterprise Premium, enables organizations to mature their AI security posture at their own pace. For security and infrastructure leaders looking to get ahead of the AI data security challenge, Chrome Enterprise is a solution worth evaluating.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.