


How telemetry brings **web defense** into the next frontier

A strategic framework for
modern enterprise defense



The browser as an AI co-worker

The browser has evolved from a passive search tool into an active AI co-worker and the primary operating environment for AI-powered and agentic work. In this new paradigm, data loss and agentic risk have emerged as primary security operations concerns for modern, AI-driven enterprises. As knowledge workers spend over half (56%) of their workday in the browser,¹ it serves as the essential gateway through which employees and autonomous agents interact with large language models (LLMs). To navigate this fundamental shift in how work is performed, organizations must establish a robust browser strategy as the strategic response to these agentic-era risks, recognizing the browser as the critical telemetry and enforcement hub for future enterprise defense.

Rather than relying on traditional endpoint security, organizations must treat the browser as a proactive telemetry engine and a centralized enforcement point that defines the future of enterprise defense. Traditional security architectures are fundamentally blind to how AI agents operate within the browser, failing to capture page-level interactions. By positioning browser telemetry as a strategic, real-time control point, security operations teams can transform browser security from a passive visibility tool into a proactive defense mechanism that continuously mitigates agentic-era risks before threats manifest into incidents.

56% of knowledge workers' days are spent in the browser¹

¹Omdia, [Securing AI Usage Starts in the Browser](#), research spotlight commissioned by Google (April 2026).



Addressing shadow AI and emerging risks

As organizations rapidly adopt generative AI and autonomous workflows, managing agentic risk and “Shadow AI”—the unsanctioned use of public LLMs where sensitive corporate data is exposed via user prompts or automated agent actions—has become paramount. Omdia reports that more than 90% of organizations allow employees to use public GenAI applications, and 92% of survey respondents express deep anxiety regarding data leakage via these unsanctioned AI applications or browsers.² Traditional network and endpoint detection and response (EDR) strategies are inherently insufficient in this landscape, as they lack the granular visibility required to inspect or control in-browser prompt dynamics and agent interactions in real time.

Because traditional EDR and perimeter tools operate outside the browser workspace, they miss critical threat vectors unique to AI-based and agentic work. Omdia research highlights that among organizations experiencing browser security incidents, data loss or leakage was the second most common attack type, cited by 38% of respondents.²

38%


of Omdia respondents cite data loss as the second most common attack type²

² Omdia, [Securing AI Usage Starts in the Browser](#), research spotlight commissioned by Google (May 2026).

³ Mandiant, [GTIG, AI Threat Tracker](#) (Google Cloud, 2026).

Threat actors increasingly exploit this gap through living off the land (LOTL) tactics, malicious extensions, and session hijacking. Crucially, without direct insight into how AI assistants interact with the Document Object Model (DOM), sensitive files, and Chrome DevTools Protocol, legacy security layers remain blind to real-time data exfiltration and unauthorized agent manipulation. Establishing a proactive browser strategy provides the centralized enforcement needed to address these agentic-era risks.

Further emphasizing this necessity, the Mandiant research confirms that state-sponsored and financially motivated threat actors are increasingly integrating AI to accelerate the attack lifecycle.³ As adversaries leverage AI as a force multiplier, the browser serves as the essential frontline for detecting and mitigating these evolving threat vectors.



Chrome Enterprise Premium as a telemetry engine

Chrome Enterprise Premium addresses these challenges by serving as a high-fidelity telemetry engine that provides deep visibility into user activity and security events. Unlike traditional tools that rely on external observation, Chrome Enterprise Premium generates real-time signals directly from the point of interaction. Threat detection and defensibility capabilities include Chrome Event Logs, investigation tools for log analytics, and connectors that seamlessly integrate alerts with a variety of tools, including Google SecOps, Splunk, and CrowdStrike.



Core Telemetry Capabilities

The urgency for these capabilities is reflected in market priorities: 59% of organizations now identify GenAI application security as the most important use case for a secure browsing solution, ranking it ahead of general data loss prevention (51%) and general web security enhancements (42%).⁴

- **Real-time signals:**
Chrome Enterprise Premium provides continuous telemetry for network events, high-risk user and agent behaviors, and suspicious domain access.
- **Chrome Event Logs:**
Comprehensive logs capture a wide array of activities, including file uploads, downloads, and sensitive data handling.
- **Evidence Locker and data snippets:**
These key features allow security teams to capture and preserve records of specific user actions, including data snippets that violate DLP policies, providing a crucial trail for forensic analysis and storing detected malware.
- **GenAI and SaaS app report:**
This new reporting capability helps IT administrators discover and govern sanctioned versus unsanctioned AI tools, identifying Shadow AI risks across the fleet.
- **Managed and unmanaged support:**
Chrome Enterprise Premium delivers a consistent control panel and telemetry stream across both corporate-managed devices and unmanaged bring-your-own-device (BYOD) environments, closing a major visibility gap for the extended workforce.
- **Integrated browser security:**
This approach offers a significant advantage for securing the extended workforce, including contractors, third parties, and BYOD users. By delivering security at the browser level with a simple user sign in, without requiring additional software or proxy, organizations can reduce costs while improving speed and resiliency at scale.
- **Advanced extension telemetry:**
Chrome Enterprise Premium also allows security operations teams to detect and respond to anomalous agent activity by surfacing detailed telemetry from browser extensions. This augments traditional endpoint detection and response (EDR) by providing granular visibility into activities EDR might miss, such as side-loaded extensions, extension-to-domain communications, and automated Google Safe Browsing protections.

⁴ Omdia, [Securing AI Usage Starts in the Browser](#), research spotlight commissioned by Google (May 2026).




Detecting agentic activity and addressing emerging AI risks

As organizations rush to embrace AI, a new backdoor has opened via AI-powered browser extensions, creating a massive, unmonitored frontier for data exfiltration and credential theft. Traditional security tools are often blind to this activity, failing to see when an extension reads sensitive financial data or writes a malicious keylogger onto a sign-in page. Chrome Enterprise Premium addresses this by surfacing advanced, page-level extension telemetry to track when AI extensions or agents interact with the Document Object Model (DOM) or inject unauthorized scripts on sensitive corporate pages. This eliminates extension blind spots where traditional EDR and network tools are blind. Furthermore, feeding this real-time telemetry into security information and event management (SIEM) and security orchestration, automation, and response (SOAR) platforms like Google SecOps enables automated playbooks to instantly uninstall malicious extensions or revoke compromised sessions, transforming browser alerts into immediate, automated defense actions.

59%

prioritize GenAI application security as the top secure browsing use case⁵

⁵ Omdia, [Securing AI Usage Starts in the Browser](#), research spotlight commissioned by Google (May 2026).



ClickFix attacks: A new frontier for agentic security

Defending from ClickFix attacks is increasingly essential for securing agentic AI environments, as these attacks manipulate users into executing system-level commands that can severely compromise browser sessions. Chrome Enterprise Premium's granular page-level telemetry serves as the primary defense to detect these deceptive interactions before they can derail sensitive, agentic workflows.

Google Threat Intelligence: Supporting agentic defense




The effectiveness of a defense strategy is tied to the quality and scale of the intelligence that fuels it. Google operates an extensive threat observatory by aggregating signals across its ecosystem, including Chrome, Android, Gmail, YouTube, VirusTotal, and Mandiant. With Google telemetry spanning billions of protected devices, organizations can maintain defensive visibility alongside emerging AI usage to drive a proactive security strategy.⁷

This telemetry is operationalized by propagating intelligence across Google's enterprise security stack. When a new threat is identified through Mandiant incident response or blocked within Chrome, that data is shared to help protect other enterprise users. This synchronization helps security leaders benefit from global insights. The system is supported by specialized threat experts across 30 countries who provide technical context to track sophisticated actors.

The business impact of integrated intelligence is measurable. According to research by IDC, organizations leveraging Google's threat intelligence have seen increases in proactive threat identification and faster mean time to detect (MTTD).⁸ These capabilities can lead to operational efficiencies, including reductions in annual costs and downtime associated with security incidents. By grounding security operations in large-scale intelligence, enterprises can maintain a more proactive security stance.

With flowing high-fidelity Chrome Enterprise Premium data directly into Google SecOps, organizations achieve visibility and proactive data loss prevention. Google SecOps serves as a cloud-native SIEM and SOAR destination, providing enriched telemetry data, a built-in detection language for rules and search, and integrated Gemini-AI for advanced security analytics.

This integration enables the transition from manual investigation to automated detection and response. For instance, security teams can build playbooks that automatically trigger responses based on browser-level alerts, such as moving a high-risk user into a restricted access group or globally blocking a newly identified malicious domain. By enriching SIEM data with granular browser telemetry, security operations teams can correlate activity across the entire user journey, identifying patterns that traditional EDR might miss, such as a user navigating to a phishing site immediately followed by an anomalous data download.



⁷ Google Safe Browsing

⁸ IDC, [The Business Value of Google Threat Intelligence](#), commissioned by Google (January 2026)

Real-world impact: Proactive visibility and action

Shifting to a proactive stance transforms how organizations handle common threat vectors. Instead of reviewing these as post-incident case studies, telemetry allows for real-time intervention and policy enforcement for a wide range of threats, not only agentic risks.



Proactive threat mitigation

Threat vector	Reactive review (traditional)	Proactive visibility (Chrome Enterprise Premium)
QR phishing	Analyzing how a QR scan led to a credential harvest and financial loss.	Real-time URL risk evaluation assesses the linked site at the moment of access, blocking the connection before the user can enter credentials.
Insider threats	Discovering data theft after a disgruntled employee has already resigned.	Monitoring the volume and speed of data access in real-time allows admins to detect and block unusual exfiltration patterns as they occur.
Malicious extensions	Investigating suspicious domain requests from an extension that has already been deleted.	Pre-deployment risk assessments for every extension and continuous monitoring of extension behavior allow for the automated removal of risky add-ons.
Mobile data loss	Limited ability to prevent data leakage on personal mobile devices.	Mobile download protection for Android and iOS blocks sensitive data downloads on unmanaged devices.
AI Agent Abuse	Manual review of audit logs after a data leakage event occurs.	Real-time detection of unusual exfiltration patterns from AI assistants interacting with the browser DOM.

Real-world threat scenarios and Mandiant insights

These real-world examples are included to demonstrate the practical application of the browser telemetry and proactive defense strategies discussed throughout this whitepaper. Discovered through Mandiant's work with customers, these case studies prove the tangible value of granular browser visibility in detecting threats—such as social engineering, unauthorized remote access, and credential harvesting—that often bypass traditional security layers, effectively bridging the gap between theoretical defense and real-world security operations. Mandiant uncovered these incidents through their managed cybersecurity services that integrate frontline threat intelligence, active threat hunting, 24/7 Managed Detection & Response (MDR), and rapid incident response directly with automated browser telemetry. By pairing human expertise with automated telemetry, these services enable security teams to swiftly detect, investigate, and neutralize advanced threats. Furthermore, these scenarios directly reflect broader industry findings from Mandiant, which highlights a notable shift toward interactive social engineering tactics and rapid threat actor handoffs.⁹

⁹ Mandiant, [GTIG AI Threat Tracker](#) (Google Cloud, 2026)

Case one

RMM software download

Summary: In a real-world scenario investigated by Mandiant, a user attempted to download 'ScreenConnect.ClientSetup.exe' from a newly registered domain. Traditional network and endpoint tools often fail to flag legitimate remote monitoring and management (RMM) executables because the binaries themselves are signed and benign. However, Chrome Enterprise Premium captures the critical context—that the download originated from a high-risk, newly registered domain. By streaming this granular page-level telemetry directly to Google SecOps, security teams can correlate origin signals, detect social engineering or unauthorized remote access attempts in real time, and proactively block or isolate the threat before execution.

Learnings

Downloading RMM software from a newly registered domain is a strong indicator of potential social engineering or unauthorized remote access. Security teams can leverage "Download Event From Newly Created Domain" and "GTI Medium Severity File Download" signatures to detect these threats proactively, often before other vendor detections trigger.



Case two

Credential harvesting attempt

Summary: Mandiant response data highlights an incident where a user navigated to a deceptive phishing URL crafted to imitate a Microsoft sign-in page. Conventional network controls and email filters can miss newly generated phishing domains or obfuscated redirects that bypass perimeter defenses. Chrome Enterprise Premium bridges this visibility gap by evaluating URL risks and navigation parameters at the precise point of interaction within the browser session. Capturing page-level telemetry allows security operations teams to instantly identify phishing navigation events, correlate suspicious URL parameters, and trigger automated playbooks to disrupt credential harvesting before users submit sensitive credentials.

Learnings

When a navigation event contains suspicious parameters, it often signals an attempted credential harvesting attack. By using “Phishing Navigation Event” signatures, organizations can identify these threats at the point of navigation, providing an opportunity to intervene before credentials are submitted.

Case three

Malicious advertisement (Malvertising)

Summary: A Mandiant investigation revealed a case where a user clicked a malicious ad, triggering an immediate high-risk executable download ‘KitchenCanvas_528034.exe’. Malvertising campaigns frequently bypass traditional EDR and network security because the initial user action occurs within standard web browsing traffic across legitimate ad networks. Chrome Enterprise Premium delivers continuous, in-browser visibility that correlates web activity with Google Threat Intelligence (GTI). This rich, real-time telemetry enables security teams to flag malicious downloads instantaneously, providing the precise evidence needed to contain threats before actors establish hands-on-keyboard persistence.

Learnings

This case demonstrates the power of integrating browser telemetry with broader threat intelligence. While the user was tricked by a deceptive advertisement, the combination of Google Threat Intelligence (GTI) and browser logs allowed for immediate identification. The incident highlights how automated alerts can ensure a threat is contained before a threat actor gains hands-on-keyboard access.

Case four

Phishing bypass


Summary: In this scenario analyzed by Mandiant, phishing emails successfully bypassed standard email security gateways and user security warnings, directing victims to malicious landing pages. Threat actors regularly utilize CAPTCHAs, proxy evasion, and cloaking techniques to prevent automated network crawlers and standard security tools from inspecting landing pages. Chrome Enterprise Premium provides essential protection at the final stage of the attack vector by capturing browser-level events such as ‘SafeBrowsing Warning Bypassed’ or suspicious DOM interactions. This granular telemetry equips security operations teams to identify bypassed threats in real time and automatically enforce remediation policies when legacy layers fall short.

Learnings

Attackers are increasingly using techniques like CloudFlare CAPTCHAs to evade detection and force users through complex landing pages. This scenario underscores the necessity of monitoring for ‘SafeBrowsing Warning Bypassed’ events. It also highlights that even when standard email protections are bypassed, browser-level telemetry provides critical visibility into the final stage of the attack—traversing to the malicious landing page—allowing teams to take action even when other vendor alerts are absent.

Essential security recommendations

To fully leverage the power of browser-based security, organizations should follow a structured, eight-step roadmap to harden their defenses and enhance visibility.



01

Centralize AI governance: Establish clear policies for sanctioned AI tools and use technical controls to monitor and block prompts containing sensitive information.

02

Enforcement in the browser: Implement access and data protections policies right in the browser, securing user and agent activity based on a managed browser experience.

03

Contextualize and enrich: Ingest browser telemetry into your SIEM to correlate signals with other security data, providing a holistic view of potential threats.

04

Automate and orchestrate: Utilize SOAR capabilities to create automated playbooks for swift remediation of browser-based attacks.

05

Technical controls: Apply technical hardening, such as granular extension controls, patch management, and URL filtering, to minimize the attack surface.

06

User awareness and training: Educate users on recognizing sophisticated tactics like QR phishing and maintaining safe browsing habits.

07

Incident Response: Develop robust plans focused on the rapid containment of browser-based threats and deep forensic analysis using captured telemetry.

08

Stay Ahead with Gemini: Leverage Gemini-powered summaries in the Google Admin console to stay informed on critical security updates, new policies, and deprecations.


Conclusion



Key takeaways for security leaders



The browser is a primary telemetry source for the AI era, serving as the essential interface where over half of the modern workday and many generative AI interactions occur.



Granular visibility is critical for securing generative AI and mitigating shadow AI risks, as 92% of organizations express concern over data leakage via unsanctioned AI applications.¹⁰



Integrating high-fidelity browser signals with Google SecOps empowers teams to correlate user activity across the entire journey and transition from manual investigation to automated response.

Global telemetry scale¹¹



5B

protected devices



3M

Users warnings issued



10B+

URLs analyzed daily



Chrome Enterprise Premium closes the visibility gaps left by traditional EDR tools, which often miss browser-based living off the land tactics, malicious extensions, and data exfiltration patterns. It offers a critical enforcement point for secure access and data protections to keep your organisation and its data safe in the age of agentic AI.

¹⁰ Omdia, [Securing AI Usage Starts in the Browser](#), research spotlight commissioned by Google (April 2026).

¹¹ [Google Safe Browsing](#)

The evolving threat landscape requires a fundamental shift in how we approach enterprise security. Organizations can no longer afford to leave the browser as a blind spot in their defense strategy. By putting the browser at the center of the security architecture and leveraging it as a primary telemetry source, security leaders can achieve the proactive visibility necessary to defend against modern, sophisticated attacks, including new agent-based threats. Closing the security gap starts with recognizing that the browser is not just an interface for work—it is the frontline of enterprise defense for today and the future.



Additional Resources

For additional strategic intelligence on adversary tactics, initial infection vectors, and emerging threat landscape trends, consult the [Mandiant M-Trends 2026 report](#).



Contact an Expert

For more information on implementing proactive browser security, contact us today.

[Learn more →](#)