IDC

SPOTLIGHT

Sponsored by: Menlo Security

The enterprise browser is now the core platform for human and AI workflows, requiring unified, session-based security. AI agents redefine access, demanding consistent policy enforcement and observability across all user types.

From Browsers as a Risk to Browsers as Security Control Planes: Securing the Next Billion AI Agent Users

March 2026

Written by: Mark Child, Associate Research Director

Introduction: The browser's evolution and the rise of agentic AI

The enterprise browser has undergone a radical transformation. Once a mere productivity tool, it is now the de facto operating system of the modern workplace. Over 85% of enterprise workflows are browser-based, making the browser both the primary attack surface and a critical enforcement point for security policy. This shift is driven by the proliferation of SaaS, the persistence of hybrid work, and, most disruptively, the integration of AI agents directly into browser and workflow contexts. As organizations move toward an agentic operating model in which AI agents act as users, the definitions of "user" and "access" are fundamentally changing. The next billion users will be AI agents, not humans. This changing dynamic demands a re-examination of browser security architectures and a new approach to securing human and non-human (agentic) identities operating at machine speed.

Historical context: The browser security journey

In the early days, enterprise security was perimeter-centric. Firewalls, proxies, and VPNs provided coarse-grained access control, with little browser-specific visibility. Threats such as drive-by downloads, cross-site scripting, basic phishing, and credential theft were relatively unsophisticated. Controls focused on endpoint antivirus, URL filtering, and basic browser hardening (e.g., disabling plugins). The browser was not yet recognized as a primary attack surface.

The rise of SaaS and BYOD blurred network boundaries. Browsers became the primary interface for business applications, and the attack surface shifted accordingly. New threats emerged, including OAuth abuse, session hijacking, malicious extensions, and shadow IT. Security controls evolved to include CASB, SSO, MFA, and browser management via MDM/EMM. However, visibility and control remained fragmented, and the browser's centrality to enterprise risk was only beginning to be understood.

AT A GLANCE

WHAT'S IMPORTANT

The enterprise browser has evolved into the core operating system for both human and AI-driven workflows, making it the primary attack surface and enforcement point for security. As AI agents proliferate, organizations must adopt session-centric, identity-aware browser security to manage risk, compliance, and operational complexity at machine speed.

KEY TAKEAWAYS

- The browser is now the universal adapter for enterprise workflows, not just a productivity tool.
- AI agents are redefining "user" and "access," introducing new risks and requiring unified, session-based controls.
- Security architectures must shift from endpoint-centric to browser-centric, ensuring consistent policy enforcement and observability for both human and non-human actors.

Zero trust principles ("never trust, always verify") and hybrid work models drove demand for continuous, context-aware security. Threats became more sophisticated, including advanced phishing, token theft, lateral movement, and supply chain attacks. Controls expanded to incorporate EDR, secure web gateways (SWGs), and the emergence of dedicated enterprise browsers and remote browser isolation. The browser was now recognized as both a critical productivity tool and a high-value target.

Modern secure enterprise browsers embed security controls and telemetry directly in the browser; enforce zero trust with continuous, risk-adaptive authorization; and offer unified visibility across user, device, and application activity. They integrate data-centric controls such as DLP and watermarking and treat identity as the new perimeter with continuous authentication. However, while it promises automation and scalability, the agentic model introduces new risks (e.g., identity sprawl, privilege escalation, prompt injection attacks, data exfiltration, and limited observability), with many agents lacking adequate monitoring and audit trails, complicating incident response and compliance.

The fundamental problem is that traditional enterprise security controls were built around human behavior, relying on user training, warnings, and endpoint posture checks. These approaches assume users act with intuition, can be educated to avoid risky actions, and are subject to controls such as two-factor authentication and behavioral analysis. However, AI agents fundamentally break these assumptions. Agents operate autonomously and at scale, executing instructions without human judgment or hesitation. They often require persistent service accounts, bypassing behavioral analytics and lacking support for human-centric controls (e.g., 2FA). Many agents interact with enterprise resources through browser sessions, especially when APIs are unavailable or incomplete, and they frequently possess broad access privileges without the benefit of fine-grained RBAC or operational constraints. This makes them susceptible to new attack vectors, including prompt poisoning and adversarial manipulation, and increases the risk of rogue or unsafe operations.

As a result, the emergence of agentic workflows introduces a new set of requirements for enterprise security. Organizations now need a unified policy framework that governs both human and agent access, with session-based controls enforced directly at the browser or web layer. Security architectures must provide consistent coverage across managed, unmanaged, and BYOD devices, ensuring that all sessions, regardless of origin, are subject to the same standards. Strong observability is essential, with real-time visibility into session activity, data movement, and agent actions. Threat prevention must extend to zero-day attacks, while file and data security controls (such as DLP and DDR) are necessary to prevent exfiltration and enforce compliance. Agent authorization must be dynamic and context-aware, reflecting the unique risk profile of non-human actors.

Regulatory and operational pressures further drive these requirements. Zero trust mandates now require continuous, context-aware access control for all identities, human and non-human. Data sovereignty concerns demand that AI workloads respect residency and privacy requirements, particularly as agents operate across cloud and edge environments. Boards and regulators increasingly expect auditability and explainability, insisting on traceability for all agent actions and decisions.

The net result is a landscape where the browser and web protocols have become the new control point for enterprise security. Nearly all work, whether performed by humans or agents, now transits the browser, making it the universal adapter for enterprise workflows, especially where APIs are lacking or incomplete. As agents deliver automation and scale, they also introduce new risks such as fraud, the rapid emergence of vulnerabilities, data loss, and the potential for unsafe actions at machine speed. When compromised, agents can inflict unprecedented damage, underscoring the need for robust, session-centric controls at the browser layer. In this context, securing the browser is no longer optional; it is foundational to managing risk in the age of agentic AI.

Definitions

When discussing endpoint protection, there are a few core terms to keep in mind. APIs let different software systems connect and automate tasks. AI is used for tasks that typically require human intelligence, such as learning and decision-making. DLP helps prevent sensitive data from being leaked or accessed without permission, while DDR tools spot and react to data threats in real time. ZTNA means every user and device must prove their identity before gaining access, no matter where they are. BYOD is simply allowing employees to use their own devices for work, which introduces additional security challenges.

On the operations side, the SOC is the team that monitors and responds to security incidents. SIEM platforms pull together and analyze security data to help spot threats and support compliance. SOAR tools automate and coordinate security responses across different systems. SASE brings networking and security functions together into a single cloud service, while SSE focuses just on the security side of that model. EPPs protect endpoint devices from malware and other threats, and EDR tools continuously monitor endpoints to detect and respond to advanced attacks.

Benefits

Transitioning from traditional endpoint security guardrails to a session-based security model fundamentally redefines how organizations approach protection, especially as both human users and AI agents increasingly drive enterprise workflows. Rather than focusing solely on the device or endpoint, this new paradigm centers security enforcement on what occurs within each session, regardless of whether the activity originates from a person or an automated agent.

The new approach: From endpoint-centric to session-centric security

In this model, the session itself becomes the primary locus for risk assessment, policy enforcement, and auditability. This shift underpins the concept of "unified browser security," a framework that supports both human and agent users. While the protocols remain consistent, the framework is designed to recognize and adapt to the distinct risk profiles of each actor type. Key capabilities of this session-centric approach include:

- **Session-based controls:** Dynamic enforcement of policies such as copy/upload restrictions, extension management, and AI interaction monitoring at runtime
- **Agent and extension governance:** Detection and control over browser extensions, plug-ins, and embedded AI agents to prevent unauthorized or risky behaviors
- **Shadow IT discovery:** Identification of unsanctioned SaaS applications and risky integrations that may bypass official security controls
- **User-centric protection:** Implementation of behavioral analytics, phishing-resistant authentication, and social engineering detection directly at the browser layer
- **Integration with XDR/SIEM:** Real-time browser telemetry that feeds into broader security operations, enabling unified threat detection and response across the enterprise

Although this approach may appear to abstract away from traditional endpoint protections, it actually enables a more fundamental and proactive security posture. Instead of merely detecting compromised access, session-centric security

focuses on identifying actions and behaviors within the session that could lead to adverse outcomes, providing a more adaptive and comprehensive defense against both human and agent-driven activity.

Trends

Over the past five years, IT has shifted from traditional algorithmic computing to AI-enabled systems, fundamentally changing how business processes are managed. With the rise of agentic AI — AI that can independently initiate actions without human input — we are seeing the emergence of a new user class. These AI agents can act much like humans in certain contexts, meaning security teams now need to protect not just people but also a rapidly growing population of autonomous agents. This shift is significantly increasing the complexity for CIOs and CISOs, who must adapt their strategies to secure both human and non-human users.

Complicating the task for CISOs is that all this is happening in an already challenging security environment. Organizations are juggling a mix of legacy and modern applications with varying levels of API maturity, supporting persistent remote work, and managing the risks of BYOD and unmanaged device access. Tool sprawl is common, with multiple access control and security products in play, leading to visibility gaps across endpoints, browsers, web gateways, and SaaS platforms. Security teams must also support flexible deployments across on-premises, private, and public cloud environments while maintaining low friction for users to preserve ROI. Additionally, workflows often require seamless collaboration between humans and agents, with agents able to request human assistance as needed.

Given these realities, the role of the browser is evolving. Instead of simply containing human-driven threats, the secure enterprise browser is becoming the central control plane for managing risks from both humans and AI agents. As agentic AI becomes more prevalent, security architectures must treat every agent as a privileged identity, enforce policies in real time, and provide unified visibility across all activity. The organizations that will thrive are those that design for secure speed, low friction, and continuous governance, making security the foundation for innovation, not a barrier.

Considering Menlo Security: Securing the browser for humans and AI agents

Menlo Security's approach to browser security is embodied by its Browser Security Platform (BSP), enabled by the Menlo Agent Runtime Security (MARS) engine, which seeks to provide session-based protection for both human users and AI agents. The BSP is built on the recognition that browsers and web protocols have become the primary execution layer for enterprise workflows, making the browser the central control plane for productivity, risk management, and security enforcement. As organizations standardize their use of browser-based applications and as AI agents become operationally critical, MARS is engineered to secure browser sessions, whether they originate from a human or an agent, independent of device or network context.

The solution's platform-centric approach is intended to address two key challenges:

- **Data starvation:** MARS gives AI agents secure, policy-governed access to necessary data, without exposing sensitive enterprise information or breaching compliance boundaries.
- **Insider threats:** Recognizing the speed and scale at which agents operate, BSP enforces strict governance, observability, and control over all human- and agent-initiated browser activity, preventing inadvertent or malicious data exfiltration and unauthorized actions.

BSP creates a secure browser runtime environment that protects both humans and agents. For document handling, it deconstructs documents into their constituent elements and reconstructs a pristine, sanitized version so that safe and compliant content is sent to users and agents. This process helps eliminate embedded threats and preserve the native user experience and document formats.

The platform is designed to operate at a global scale, supporting millions of concurrent browser sessions. Its egress infrastructure is engineered to avoid automation blocks such as captchas, and programmatic geolocation enables access to region-specific content and government websites. All web and file content is processed within the runtime, removing threats so only compliant data is ingested. Data loss prevention and data detection and response controls are applied before and after AI processing to enforce enterprise data policies and prevent leakage. Every agent transaction is governed, logged, and auditable, supporting compliance and data governance requirements. MARS also provides an access framework for zero trust network access (ZTNA), allowing cloud-based agents to securely access on-premises enterprise applications without exposing sensitive infrastructure.

MARS's architecture supports secure connectivity and access to modern SaaS, legacy web applications, and internet resources. It operates seamlessly across managed, unmanaged, and BYOD endpoints. Typically consumed as a cloud service, it can also be deployed on premises or in private cloud environments. Consistent security and policy enforcement is provided for both human and agent users, regardless of connection method or location. To counter advanced threats, MARS offers zero-day protection against phishing, fraud, and malicious web content at the session layer. Data security is enforced through granular controls over sensitive data access, exfiltration pathways, and browser-based leakage prevention.

A unified policy engine and management experience underpin the platform, allowing security teams to define, enforce, and monitor policies across all sessions and user types. Telemetry from every session feeds into SOC workflows, compliance and audit processes, and incident response pipelines, integrating with SIEM and SOAR platforms for comprehensive observability and auditability. Menlo's reference architecture visualizes both human and agent browser sessions under a single security enforcement layer, with all activity logged and available for real-time or retrospective analysis (see Figure 1).

Figure 1

Menlo's browser security platform

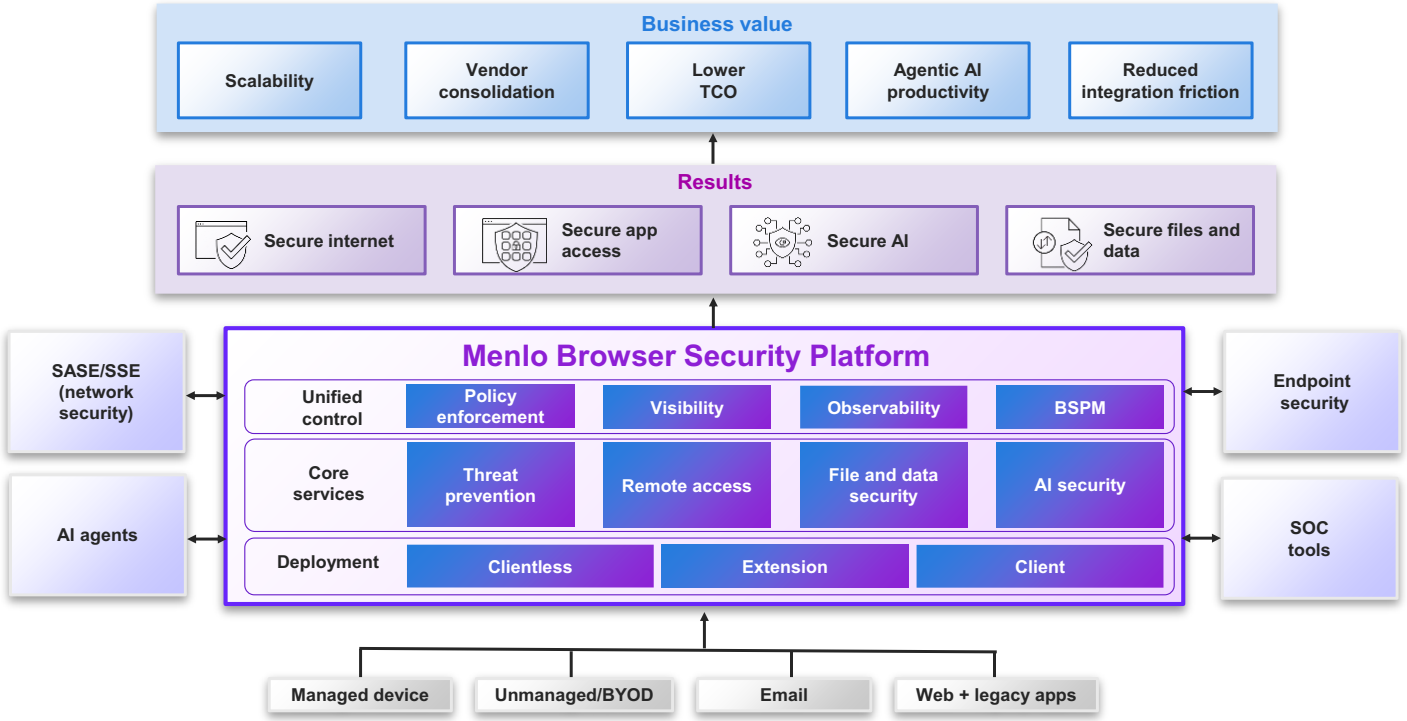

Source: Menlo Security, 2026

This platform-centric approach empowers organizations to safely enable agentic AI, allowing secure access to AI tools, embedded AI applications, and enterprise data while preventing unsafe agent actions. The platform is also designed to support secure third-party and unmanaged access, handle phishing and fraud, defend against malware and ransomware via files, and enable robust data loss prevention. As AI agents proliferate, MARS's strategy treats every agent as a privileged identity, enforces policy in real time, and provides unified observability across all browser and agent activity, making security the foundation for innovation in the hybrid human/agent enterprise.

Challenges

Despite its unified approach, Menlo faces several challenges in the evolving security landscape. The convergence of browser security with broader categories such as SASE, SSE, EPP, and EDR creates confusion for organizations that must navigate overlapping feature sets and integration claims. Large SSE vendors are increasingly bundling browser controls, intensifying competitive pressure and raising the bar for differentiation. Another challenge is the heterogeneity of agent implementations. Not all agents interact with enterprise resources via browser sessions; some use APIs directly, which may bypass browser-centric controls. This diversity requires Menlo to maintain flexible enforcement mechanisms and clear governance models for defining and managing agent identities, as well as robust audit trails for autonomous actions. Finally, Menlo's platform claims (e.g., true consolidation, a unified policy engine, and operational simplicity at scale) must be substantiated in real-world deployments. Organizations will scrutinize whether Menlo can deliver on these promises while supporting the rapid growth of agentic workflows and maintaining ease of management across complex, distributed environments.

Conclusion: Browser security as the foundation for the agentic enterprise

Agentic AI is expanding the number and nature of users in the enterprise. Browser and web protocols remain the dominant interaction layer for humans and agents. Therefore, enforcement and control must move to the browser, with policies that cover both user types. Menlo and similar platforms are positioned to help enterprises modernize access, enable agentic AI adoption, and capture ROI without increasing risk or complexity. The next billion users will be AI agents, and the browser will be their operating system and their control plane.

The next billion users will be AI agents, and the browser will be their operating system and their control plane.

About the analyst

Mark Child, *Associate Research Director*

Mark Child of IDC's Security Group leads the group's Endpoint Security and Identity & Digital Trust (IDT) research. He monitors developments in security technologies and strategies as organizations address the challenges of evolving business models, IT infrastructure, and cyberthreats. Mark's coverage includes in-depth security market studies, end-user research, white papers, and custom consulting.

MESSAGE FROM THE SPONSOR

Menlo Security is the pioneer of unified browser security, protecting the modern enterprise's dual workforce: humans and AI agents. Our Browser Security Platform is the only solution that provides a unified trust layer for the modern enterprise, delivering architectural immunity to all actors, both agents and humans.

By focusing security on where the work happens — the browser session — Menlo provides industry-best zero-day threat prevention that eliminates threats before they reach the user device or agent, advanced file and data security that keeps users safe and productive, and secure access to applications. These key capabilities are tailored for the security and access needs of users and agents within a unified control, visibility, and policy framework.

Menlo does not just bundle security features; we collapse the distance between security, productivity, and innovation. For more information, visit https://www.menlosecurity.com/product.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
One Beacon Street
Suite 33100
Boston, MA 02108, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com