




White Paper

The 2026 CISO Guide to Secure AI Browsing & GenAI Risk Management

How to confidently enable AI everywhere — without losing control.



Introduction

Generative AI has officially moved into the browser. Every page, every app, and every workflow now has a copilot, sidebar, or embedded model. The browser has become the operational plane of AI — where enterprise data is viewed, manipulated, and now acted upon.

The wins are undeniable:

- ✓ Automated reporting
- ✓ Personalized sales outreach
- ✓ Accelerated development
- ✓ Instant knowledge access

But in 2025, CISOs learned a hard truth:

**AI is not just reading your data —
it is now capable of executing actions across your enterprise.**

This creates a **new and urgent security mandate**: govern AI actions at the browser level, before they can cause harm.

A Guide Built for Modern Security Executives

This guide equips CISOs and IT Leaders to:

- ✓ Navigate the new landscape of AI-driven browsers
- ✓ Understand emerging risk categories unique to AI in the browser
- ✓ Recognize the deficiencies of consumer AI browsers now entering the enterprise
- ✓ Deploy Zero-Trust for AI actions — not just identities
- ✓ Prepare for 2026's explosive shift to agentic enterprise AI

By adopting the strategies here, organizations protect productivity and compliance while enabling AI innovation at scale.

The Evolving Landscape of AI in the Browser

In enterprises today, GenAI shows up in five browser-centric ways:

In-Browser AI Modality	Examples	New Risk	Security Need
Embedded App Copilots	Salesforce, M365 Copilot, Workday	Excessive SaaS privilege & silent permissions escalation	RBAC per action & per field
Public AI Web Apps	ChatGPT, Claude, Perplexity	IP/PII prompt leakage to external training	Prompt DLP + session policy
AI Browser Sidebars & Overlays	“Ask this page” copilots	Unmonitored cross-app data movement	Tab isolation & access segmentation
Marketplace Extensions/Plugins	Chrome Extensions, Workspace Add-ons	Off-book data extraction & local model storage	Plugin visibility + policy governance
OS-Integrated AI Browsers	Edge Copilot+, Chromium AI	Omnibox injection & auto-action flows	Action audit + token-level control

Summary

AI is spreading faster **through browsers** than through SaaS vendor roadmaps.

The browser is now the most privileged application in the enterprise.

New Application Taxonomy for 2025 AI Browsing

Borrowing structure from traditional sanctioned/tolerated/unsanctioned taxonomies, but modernized to include **action authority**:

Classification	How AI Behaves	Business Benefit	Security Threat	Required Control
Sanctioned AI Agents	Operate in-browser with enterprise policies	High	High	Zero-Trust governance of action, movement, storage
Tolerated Assistants	Limited to content generation	Medium	Medium	Guardrails around sensitive content
Shadow AI Plugins	Hidden automation implementations	High	Critical	Real-time discovery & blocking
Consumer AI Browsers	Designed for productivity, not governance	High UX	Severe	Replace or wrap with secure AI browser

Deficiencies in Consumer AI Browsers

Consumer AI Browser Claim	CISO Reality
“We apply DLP”	They only inspect text , not actions
“Model doesn’t store your data”	But prompts enter model-resonant memory
“Secure by design”	No RBAC for AI agents — everything is performed with user super-identity
“Fine-grained controls”	Controls exist at the SaaS layer, not the AI agent
“Enterprise integrations”	Each plugin introduces new attack surface with no central visibility

In short:

Consumer AI browsers were built for creativity — not safety.

Top AI Browser Risks Discovered in 2025

Borrowing structure from traditional sanctioned/tolerated/unsanctioned taxonomies, but modernized to include **action authority**:

Risk Vector	Description	Business Impact
Indirect Prompt Injection (IPI)	Malicious websites inject hidden commands into a copilot	Unauthorized transactions, silent data export
Cross-Site Action Leakage	Agents act across tabs without governance	CRM → Slack → Gmail → breach in seconds
Omnibox Over-Authority	AI can execute OS-level shortcuts, downloads, deletions	Ransomware delivery, credential theft
Shadow AI Plugin Ecosystems	Extensions become covert data extractors	Unrecoverable IP loss
Agent Identity Collapse	Browser copilots inherit user privilege without scope	SOX, HIPAA, GDPR violations
Model Residual Data Poisoning	Sensitive data contaminates training sets	Regulatory exposure beyond SOC reach
Invisible Clipboard Channels	Copy-paste becomes a high-volume data leak path	Hard-to-detect incremental exfiltration

What changed this year:

Attackers no longer need to steal data directly. They just convince an AI agent to move it for them.

Zero-Trust for AI Actions: The New Mandate

Security must shift from:

Yesterday		Tomorrow
Block risky sites	----->	Govern risky actions
Protect files & objects	----->	Protect context & intent
Isolate browsers	----->	Instrument browsers
SASI / API enforcement	----->	Real-time inline enforcement
Data-at-rest DLP	----->	Agent-in-motion DLP

If an AI can act, it must also be accountable.

Only a secure AI browser can govern:

- What actions an AI **can** perform
- Where data **can** flow
- Who is responsible
- How the action is **logged and auditable**

2026: The Year of Regulated AI Browsing

Regulators are already drafting enforcement:

- Prompt exposure = **data breach event**
- Automated SaaS actions = **regulated transactions**
- Agent logs = **audit artifacts**
- Copilot privilege = **identity governance**

CISOs should expect:

- ✓ Mandatory AI action auditing
- ✓ AI DLP enforced at the browser layer
- ✓ Controls on AI-initiated system changes
- ✓ Governance over model-resonant memory
- ✓ Coverage of unmanaged devices via Secure AI Browsers

Mammoth Cyber Call to Action

Your browser needs to be your most secure application.

Enable Secure AI Browsing with:

- ✓ Enterprise AI Browser with **Identity-anchored Zero-Trust**
- ✓ Dynamic **AI Action Governance**
- ✓ Browser-native **Data Loss Prevention**
- ✓ SaaS-aware **Privilege Segmentation**
- ✓ Real-time **Plugin & Agent Discovery Blocking**
- ✓ **Cross-domain tab isolation** & secure remote rendering

We allow enterprises to adopt AI browsing **confidently**, without limiting innovation.

What Success Looks Like

Metric	Leading Indicator
AI Adoption Safety	100% sanctioned AI browsing sessions
Governance Strength	0 unauthorized AI-executed actions
Productivity Gains	Reduced ticket load & time-to-deliverables
Shadow AI Reduction	Continuous drop in unapproved extensions
Compliance Assurance	Audit-ready browser action logs

Security doesn't slow AI adoption — it powers it.

Securing the AI Future — Starting Now

2026 will bring more agentic copilots, more embedded assistants, more automation — and more opportunity.

The question for CISOs:

**“Will AI browsing become your greatest accelerator —
or your largest unmonitored attack surface?”**

With Mammoth Cyber, you choose acceleration.



Mammoth Cyber is a Palo Alto-based cybersecurity company pioneering the Enterprise AI Browser—a secure control plane for how organizations interact with AI, applications, and data. Unlike consumer-grade AI browsers, Mammoth Cyber is built for enterprises that require data loss prevention, access governance, and zero-trust guardrails to safely adopt AI. With advanced controls for multi-tenant RBAC, AI-aware DLP, and trust-circle enforcement, Mammoth Cyber ensures that employees, contractors, and AI agents can work productively without exposing sensitive data.

For more information, visit mammothcyber.com or contact sales@mammothcyber.com.