Manifold launches first-of-its-kind solution to trace and govern AI agents across browser and endpoint runtime

Browser support expands Manifold's AI security platform to AI assistants and browser-driving agents, giving security teams visibility into what AI actually does

LONDON, UNITED KINGDOM, September 16, 2026 /EINPresswire.com/ -- As AI agents move to the browser, Manifold Security, the agentic security company protecting organizations from the risks created by AI agents, today announces browser support for its AI governance platform, extending its endpoint coverage to AI activity in managed Chromium-based browsers on macOS, Windows and Linux. The launch introduces a first-of-its-kind approach to tracing and governing AI agent activity across browser and endpoint runtime from a single control plane, following the agent from the initial user prompt through its tool calls to the final outcome.

The launch comes as the capabilities and risks of AI agents are becoming increasingly visible, dominating business and government agendas. OpenAI's GPT-6 Astra pushed the capabilities of AI systems further, while the July Hugging Face incident demonstrated how autonomous AI agents can perform thousands of actions and interact with external infrastructure when pursuing a goal.

More recently, researchers have identified OpenAI agents using external websites to communicate and bypass restrictions. As agents gain the ability to search the web, interact with applications, use authenticated sessions and take actions across external systems, security teams need visibility into not just which AI tools employees are using, but what those systems are actually doing.

Manifold Security's new capability covers two types of browser AI: AI assistants used in browser tabs, including claude.ai, Claude Code on the web, Claude Cowork, ChatGPT, Gemini, Cursor and Perplexity; and browser-driving agents that operate the browser on an employee's behalf, including Claude in Chrome, Replit, ChatGPT and Gemini agents. Manifold's platform is the first platform capable of monitoring Claude in Chrome.

For AI assistants in browser tabs, Manifold features:

- Prompt and site controls – Block a denied prompt before it leaves the browser and prevent a denied AI site from loading and enforce corporate account rules where required.

- Real-time policy enforcement – Flag and record policy breaches as they happen, with violations tied to the user, device and conversation.
- Activity-based detection – Detect activity based on what the AI did, including tool calls, where data went and combinations of capabilities.
- AI inventory – Include MCP servers, skills and plugins connected to browser AI in the same AI bill of materials as endpoint activity.

For browser-driving agents, Manifold provides:

- Agent activity monitoring – Applies the same policies and detections used for endpoint agent tool calls to agents operating the browser.
- Auditable activity timelines – Reconstructs the agent's activity as a timeline, including the prompt, each tool call, approval decision, result and page-level evidence, correlated by browser tab and time.

Manifold's browser support runs through the same policy engine, inventory and audit pipeline as its endpoint coverage, with no additional governance layer or separate set of policies to maintain. It has been running in customer deployments for the last several months and is available today to customers using managed Chromium-based browsers on macOS, Windows and Linux.

Customers do not need to install the extension themselves, as it can be deployed through existing mobile device management tools. The extension's permission scope is deliberately small. Host access is limited to the AI surfaces it governs, it does not read browsing history, and the only user data it reads is the signed-in account identity, used to enforce account policies. Conversation content is not captured unless an organisation explicitly enables prompt capture. Monitoring an agent that drives the browser, such as Claude in Chrome, uses the same browser access the agent itself relies on, and nothing more.

Oleksandr Yaremchuk, Co-Founder and CTO of Manifold Security, commented: "Security is changing because AI is changing. GPT-6 Astra and the incidents we've seen around autonomous agents, from the Hugging Face breach to agents interacting with external websites, are all pointing in the same direction: AI is no longer just generating an answer. It can take action, make decisions and keep going when nobody is watching.

"The browser is becoming one of the places where that happens. If an agent can operate through a user's browser, security teams need to know what it accessed, what it did, what it was allowed to do and what happened afterwards. You cannot govern an agent you cannot see. Manifold brings that visibility and control into the browser, using the same policies and audit trail as the rest of the enterprise."

Neal Swaelens, Co-Founder and CEO of Manifold Security, added: "As AI agents move from answering questions to taking actions, the browser is becoming another part of the agent

runtime. Security teams need to see what those agents actually do, not just what they were prompted to do. We're bringing that visibility and control to the browser today, with more browser platforms and AI surfaces planned throughout 2026."

Sam Milligan
Manifold Security
+44 20 7846 7860
email us here

This press release can be viewed online at: https://www.einpresswire.com/article/942444418

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.