cyberhaven

Govern AI at the *Speed* It Operates

Inventory AI apps, autonomous agents, and the data they touch, so you can adopt AI without losing control.

The Problem

Autonomous agents like Claude Code, Codex, and Copilot now run on endpoints, inherit employee identity, access production data, and act at machine speed, with no traditional security control seeing it. Legacy DLP was built for browser-based, human-speed interactions. It can't distinguish a corporate ChatGPT instance under governance from a personal account leaking IP, or detect goal hijacking, privilege abuse, and data exfiltration across multi-step agent workflows. Shadow agents and unmanaged MCP servers expand the attack surface faster than security teams can catalog it.

60% of organizations have adopted endpoint-based agentic AI apps

49.5% of developers use AI coding assistants

39.7% of AI interactions involve sensitive data

How Cyberhaven Solves It

Cyberhaven Flow, the AI-native data security platform for the agentic enterprise, combines a Data Lineage graph, AI-powered content inspection, and a lightweight endpoint agent to give security teams real-time visibility into AI apps, agents, and data flows, across SaaS, endpoints, and developer environments. Policies are enforced based on behavior and data context, not assumptions.

01 Discover **02 Monitor** **03 Score** **04 Control**

KEY OUTCOMES

- Eliminate AI Blind Spots
- Govern Agents Before They're a Liability
- Stop Data Leaks at Machine Speed
- Score Risk Automatically
- Accelerate AI Adoption
- Investigate 5x Faster

CAPABILITIES

Shadow AI Discovery
Discovers hundreds of AI tools across endpoints, browsers, CLIs, and IDEs; existing and emerging, standalone and embedded, personal and corporate accounts.

Agentic AI Discovery
Reconstructs full execution lifecycles for Claude Code, Codex, Copilot, and other agents, including tool calls, data access, API invocations, and multi-turn context.

Data Lineage for AI
Traces data origin, movement, and transformation across every AI interaction. Connects agent actions to data for full forensic context.

AI Data Flow Control
Runtime guardrails block, warn, or redact at the prompt and response level. Plain-English risk explanations replace generic block pages.

AI Usage Insights
Adoption trends, power users, sensitive data shared with AI, and protections applied from a single console.

AI Risk IQ
Scores AI apps and agents across five dimensions: data sensitivity, model integrity, compliance adherence, user access, security infrastructure.

PLATFORM COVERAGE

GenAI & Agents
ChatGPT, Copilot, Gemini, Claude, Claude Code, Codex, Perplexity, open-source frameworks, MCP servers, and 100s more

Endpoints & Dev Tools
Windows, macOS, IDEs, CLIs, GitHub, GitLab, Jira, Confluence

Cloud & SaaS
Google Drive, OneDrive, Dropbox, Slack, Teams, Salesforce, 300+ cloud apps

Why Cyberhaven

Legacy tools detect that an employee opened ChatGPT. They cannot trace what an autonomous agent read, transformed, and acted on across dozens of API calls, or tell a governed corporate instance from a personal account leaking IP. Cyberhaven was built around data, not perimeters. A Data Lineage graph tracks every piece of data across every workflow. Combined with AI-powered content inspection and a lightweight endpoint agent, this is purpose-built security for a world where agents operate as autonomous actors inside enterprise systems.

Ready to govern AI at the speed it operates?

See how Cyberhaven inventories AI apps, agents, and data flows, and the risks your current tools are missing.

Data Security for the Agentic Enterprise

Request a demo

cyberhaven.com/request-demo

Cyberhaven is a leader in Data Security for the Agentic Enterprise, trusted by the world's most innovative companies to protect their most sensitive data. AI changed work. Cyberhaven protects it. Learn more at cyberhaven.com.