cyberhaven

Agentic Data Security

See which AI apps and agents are in use, the data they touch, and the risk they create, so your teams can support AI adoption and stay in control.

Your business wants the speed and productivity of AI. Your security team needs to know what data that AI reads, changes, and sends, and who or what is acting on it. Those two goals do not have to compete.

Cyberhaven is the leader in Data Security for the Agentic Enterprise, tracing the full lifecycle of your data, adapting protection to changing context. With the Agentic Data Security product in Cyberhaven Flow, security teams discover AI use, score risk, and control risky data flows, so AI adoption moves forward with the visibility and controls the business needs.

39.7% of AI interactions involve sensitive data

509% growth in endpoint AI agents

49.5% of developers use AI coding assistants

Source: Cyberhaven Labs.

From AI Chat to Agents That Act on Data

AI use has moved beyond a browser tab. Employees now work with desktop assistants, coding agents, command-line tools, AI features embedded in everyday apps, local models, open-source agent frameworks, and Model Context Protocol (MCP) servers that connect agents to tools and data.

These agents open files, invoke tools, and complete multi-step tasks under an employee's identity, often without oversight. A log entry that says an AI app was used does not tell you what the agent read, what it did, or where the output went.

Agents also create new artifacts at machine speed. A summary, a slide, or a code file can carry sensitive context from its source material, even when the new file has no label to show it. Meanwhile, security teams often cannot tell a governed corporate AI account from a personal one, and an alert without data context, generated by an agent, is increasingly difficult to investigate and resolve.

Three Gaps Security Teams Need to Close

Agents act on sensitive data across workflows, and security teams need data context, agent visibility, and controls that keep pace.

Shadow AI gap

Security teams need a current view of the AI apps, agents, accounts, and MCP servers in use. Manually maintained catalogs fall behind as fast as new tools appear.

Workflow-context gap

Knowing which tool is installed does not reveal what data an agent accessed, what actions it took, which identity it acted under, or where the data went. Without that context, teams cannot separate routine work from real risk.

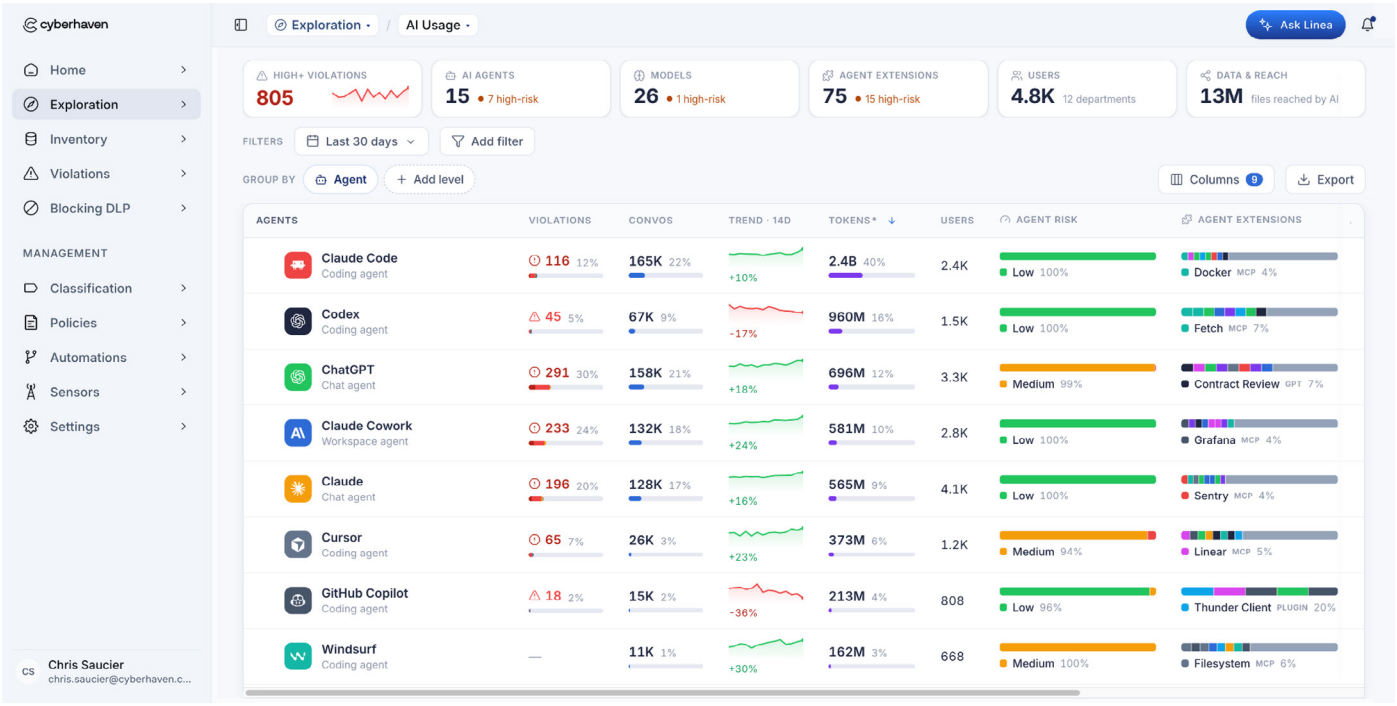
Control gap

Controls built for browsers and human users may not see endpoint agents or multi-step workflows in the same way. The challenge is coverage: protection needs to reach the places where agents act on data.

Protect Workflows, Not Just Data

Cyberhaven Flow combines data lineage, AI content inspection, and an endpoint agent to give security teams visibility into AI apps, agents, and data flows across SaaS, endpoints, and developer environments.

Lineage is the context layer. It links each piece of data to its origin, its movement, and every transformation along the way, then relates that history to the agent activity and the identity involved.



Agentic Data Security turns that context into action in four steps, and each step builds on the one before it.

Discover, Monitor, Score, Control

1. Discover AI Apps and Agents

Inventory sanctioned and unsanctioned AI apps (e.g. corporate vs. personal), agents, accounts, and MCP servers across supported endpoints, browsers, command-line interfaces (CLIs), and integrated development environments (IDEs).

2. Monitor Agent Activity and Data Lineage

See supported activity context such as tool calls, data access, and multi-turn interactions. Agentic Data Security connects each agent action to the files and data involved and the identity behind it. Teams can trace data origin, movement, and transformation across AI interactions to assess risk and investigate events.

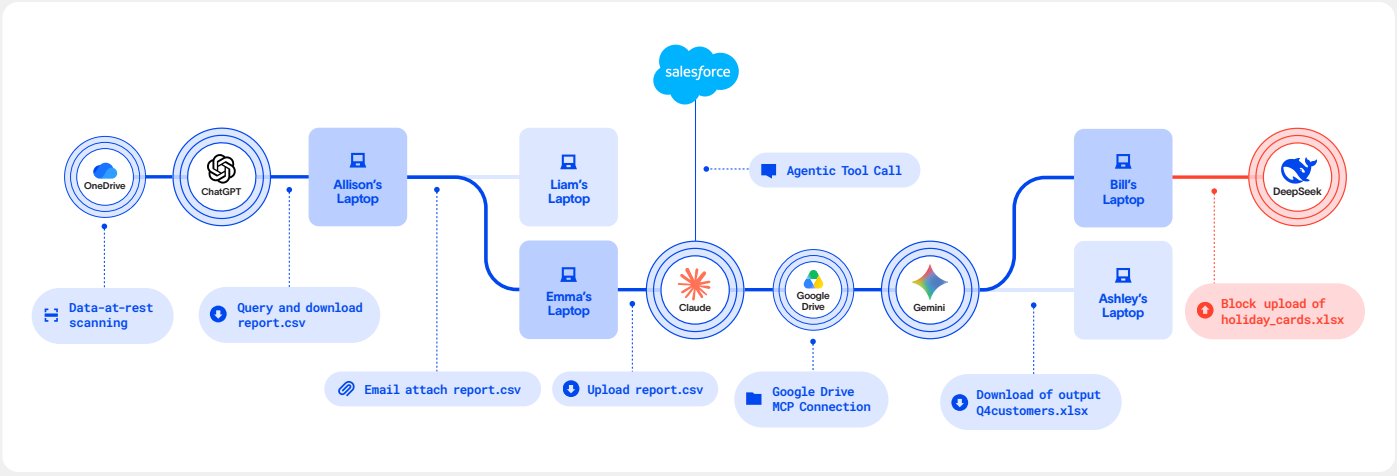
3. Score Risk with AI Risk IQ

AI Risk IQ scores AI apps and agents across five dimensions: data sensitivity, model integrity, compliance adherence, user access controls, and security infrastructure. Teams can use the scores to prioritize investigations and governance decisions. AI Usage Insights adds adoption trends, power users, sensitive data shared with AI, and the protections applied.

4. Control Risky Data Flows

Apply guardrails that block, warn, or monitor against risky agent activities. Policies consider data context and agent behavior, so teams can stop risky activity and keep legitimate AI use moving.

How It Fits Together



Context is what separates Flow from a tool inventory: lineage and content inspection give every risk score and policy decision the history of the data behind it.

Proven Data Security, Extended to Agentic Workflows

Cyberhaven Data Loss Prevention (DLP), Data Security Posture Management (DSPM), and Insider Risk Management (IRM) capabilities extend to AI agents, providing end-to-end visibility and control for data.

Agentic DLP

Stop sensitive data and credentials from leaving through sanctioned and unsanctioned AI agents. Lineage-aware policies follow data across endpoints, browsers, and the cloud, so protection holds even when an agent copies, summarizes, or fragments the original.

Agentic Data Security Posture

Discover and inventory AI agents, MCP servers, and tools, then enforce allow-lists and block shadow servers. Trace AI-generated content to its originating agent, session, and source data, and find sensitive data in past agent activity.

Agentic IRM

AI agents are the new insider risk. Agents act with the access of the people they serve, at machine speed. Investigate any agent session step by step, with data lineage, identity attribution, and sensitivity context at each step, and see agent behavior alongside the people behind it.

Outcomes for Security and the Business

Reduce AI blind spots

Know which AI apps and agents are in use across the endpoints, developer tools, and cloud services you cover.

Investigate with context

Connect agent activity to data, identity, and destination, so analysts work from more than an app name or an isolated event.

Protect sensitive data

Apply data-aware controls to risky AI flows before risk becomes loss.

Support governed AI adoption

Use usage and risk insights to decide which tools to sanction, restrict, monitor, or review.

Across the Cyberhaven Platform, Teams Have Seen

5x	2x	90%
faster incident investigation	faster incident resolution	fewer false positives

Where Flow Helps

Scenario	Security question	How Flow helps	Outcome
A developer runs an AI coding assistant with access to source code and credentials	Which agent is active, what did it access, and is any flow risky?	Discovery of coding agents, activity monitoring, and data lineage tied to the developer's identity	Risky code and credential flows are identified and controlled.
An employee uses a personal AI account while a sanctioned enterprise tenant is available	How much AI use happens outside governed accounts, and with what data?	Corporate versus personal account discovery and AI Usage Insights	Teams prioritize governance and steer users to the sanctioned tenant.
An analyst investigates a possible AI-related exposure	Was sensitive data accessed, transformed, or sent through an AI interaction?	One view connecting agent activity, data lineage, and identity	Faster, evidence-based conclusions with fewer manual lookups.

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

Full lifecycle context
Cyberhaven traces data from origin through every movement and transformation, so every AI event is judged against the full history of the data.

Agent-aware visibility
Flow brings agent behavior and data access together, so teams see what an agent did with data after an AI tool was opened.

A data security foundation
Agentic Data Security sits within Cyberhaven's broader platform, which also delivers DLP, DSPM, and IRM.

Built where work happens
An endpoint-centered approach extends across AI, developer, and cloud and SaaS surfaces, so Flow can take action when and where it matters.

> "The provenance that Cyberhaven provides is critical to building your data program for AI, or for any cybersecurity measures." — **CISO, Life Sciences Firm**

See Agentic Data Security in Your Environment

Request a demo to see the AI apps, agents, workflows, and risk in your environment.

Request a demo

cyberhaven.com/request-demo

AI Changed Work.
We Protect It.

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. Learn more at cyberhaven.com.