harmonic | Harmonic Command

Harmonic Command: Enable Safe AI Use for Humans and Agents

Visibility and Governance For AI Everywhere

The most ambitious organizations are moving fast with AI across every surface, from employee workflows to autonomous agents operating inside enterprise systems. Harmonic Command is the governance layer built for that pace.

Harmonic Command at a Glance

Harmonic Command gives security teams real-time control of human and agentic AI workflows. Our specialized small language models understand intent and context, detecting sensitive data without regex rules or manual labeling.

When a risk is identified, Harmonic intercepts the interaction before data leaves the business and guides the employee or agent toward safer behavior, all within milliseconds.

Built for Every Way Your Organization Uses AI

AI exists everywhere employees work. Developers use Claude Code. Marketing teams create assets in Canva. Business analysts run data through ChatGPT Desktop. And increasingly, autonomous agents connect to internal systems through MCP servers, acting across tools like Salesforce, Jira, and internal platforms on behalf of your teams.

Harmonic Command governs all of it within a single platform. Human and agent interactions are held to the same policies, analyzed by the same models, and managed from the same place.

Key Features

- **Prompt level visibility across 1,000+ apps**, including tools adopted without IT approval or used within agentic workflows.
- **Understand AI use cases across the business** with Usage Intelligence that maps activity to business impact.
- **Detect and prevent sensitive data leakage** with small language models that understand context, not just patterns.
- **Coach users and agents with inline interventions** that warn, nudge, or block at the point of data loss without stopping the workflow.

See Everything

Know exactly how your org leverages AI and where risk exists

Empower Your Workforce

Enable everyone to safely use the AI tools that help them the most

Govern Agents

Control what they can do and what they can share without interrupting workflows

Leverage True AI Usage Intelligence

Most organizations can tell you which AI tools employees are using. That doesn't tell you why it matters. Usage Intelligence changes that.

Harmonic automatically classifies AI interactions into business use cases, turning raw activity into strategic insight. Instead of generic tags or app-level analytics, you see how AI work actually happens across your organization, specific to your teams, your tools, and your business.

You can see which tools are delivering real productivity gains, which use cases are driving the most value, and where employees are flocking to unapproved tools because sanctioned options aren't meeting their needs.

When sensitive information is shared, Harmonic doesn't just flag it. We map it to the regulatory frameworks and compliance standards that matter to your business, including GDPR, ISO 27001, and HIPAA. You're able to tie policies directly to GRC initiatives.

This is the foundational layer of AI governance.

Before you can implement effective controls or make informed AI investments, you need to understand where AI is delivering ROI, where it's creating risk, and why employees are choosing the tools they choose. Harmonic delivers on all of it.

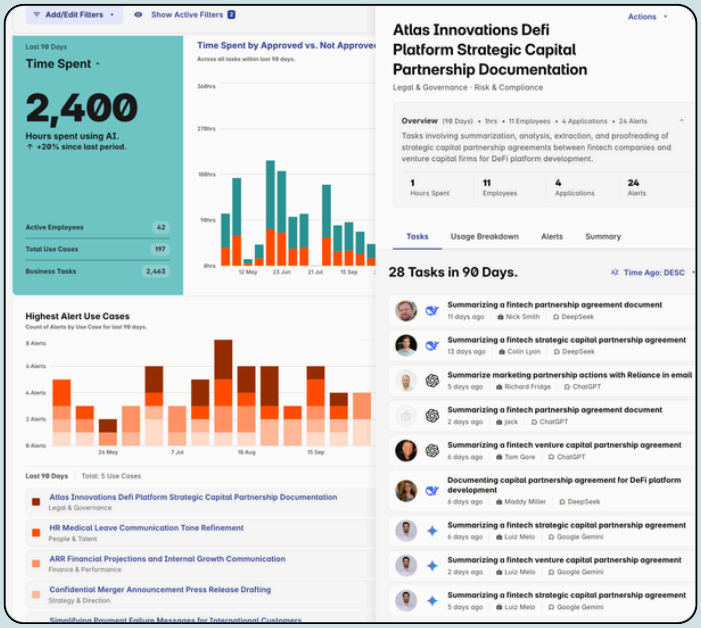

Actionable Usage Insights

Gain Deep Insight Into Every AI Application

Harmonic doesn't just tell you which apps are in use. We give you a detailed profile of each one, covering both the productivity value it delivers and the risk it introduces.

For every application, you see who's using it, how often, what use cases it supports, and what sensitive data is being shared with it. You can quickly identify which tools are driving real business value and which ones are quietly creating exposure.

Each application also carries a risk score based on factors that security teams actually care about, like whether the tool trains on your data, where the vendor is based, and what security certifications the vendor holds.

You're armed with an objective, consistent way to evaluate the hundreds of AI tools that show up across your organization.

Harmonic also provides built-in workflows for approving or flagging applications and grouping them by risk level.

You can tier your apps into approved, monitored, and unsanctioned categories, and use those groupings as the foundation for policies of what informtion users can share with each category.

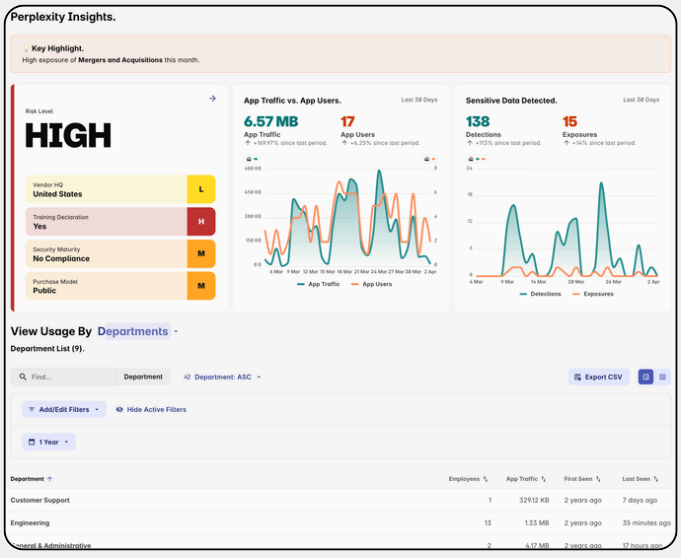

Complete Visibility by Application

Detect Sensitive Information and Coach Users in Real Time

Traditional tools look for patterns and keywords. They can tell you that something looks like a credit card number, but they can't tell you that an employee just pasted a full customer complaint record into an unsanctioned chatbot, or that a developer shared proprietary architecture details with a coding assistant that trains on user inputs. Harmonic Security can.

Our small language models read the full interaction, not isolated strings, and make human-like judgments about whether sensitive data is about to leave the business. They easily parse through complex prompts or file uploads with an understanding of prompt context and intent. Results are returned within 200 milliseconds.

When sensitive data is detected, Harmonic intercepts the prompt before it's sent and shows the employee exactly what was flagged and why.

Security teams configure how users are coached based on the data type, the application, and the user's role. Employees can be given options to cancel the input, edit it, or override with an explanation. Intervention text, colors, and branding are fully customizable.

The full context of every interaction is available in the Harmonic portal, enriched with identity provider data, so your team always knows who shared what, in which tool, and what they were trying to accomplish. The result is confident action on the interactions that matter and a culture of secure AI adoption instead of shadow AI workarounds.

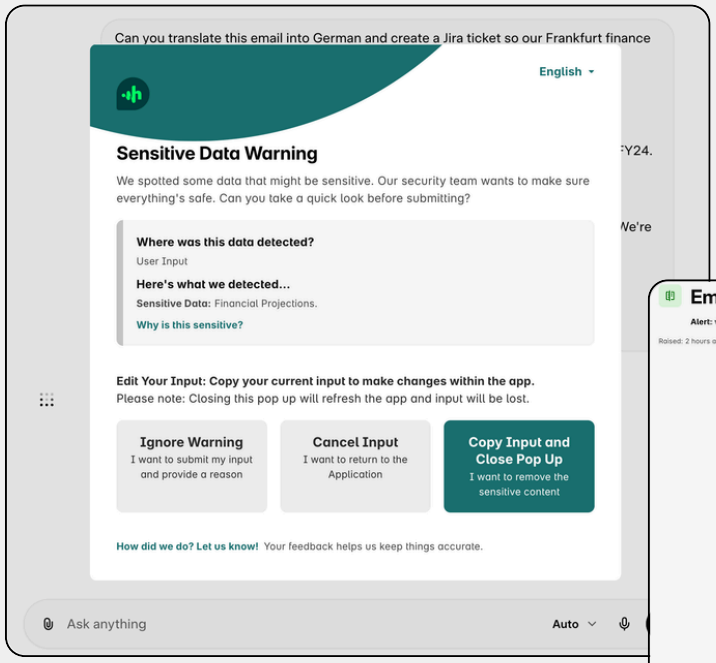

Customizable Notice Displayed to End Users

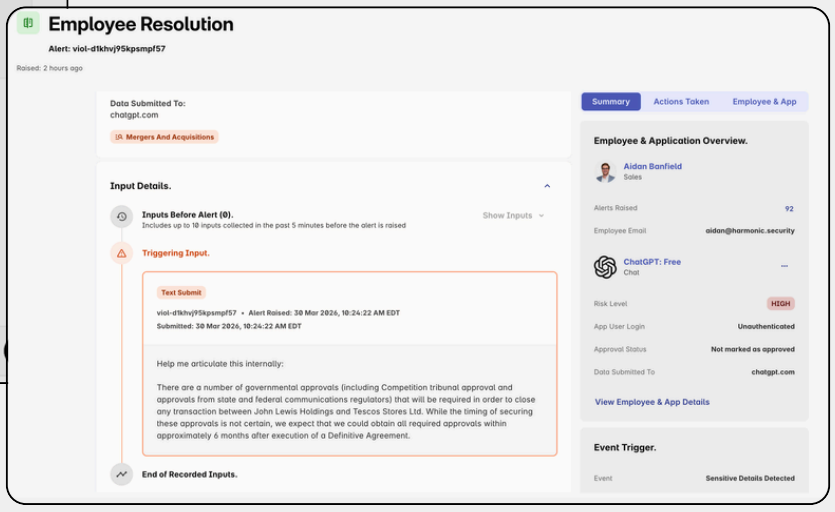

Complete Data of Each Alert and Resolution

Discover and Control Agentic AI Across Your Environment

AI agents are only as safe as the systems they connect to. Harmonic automatically discovers every MCP client and server in use across your organization, whether official vendor integrations, locally built servers, or experimental setups your developers have spun up on their own.

Once discovered, security teams get full visibility into which employees are using which clients and servers, how frequently, and what actions those agents are taking. From there, you set the rules.

Harmonic lets you define granular permissions at the agent, server, and action level. You decide which systems an agent can read from, which it can write to, and which destructive actions like deletes or overwrites are off limits.

Agents keep the access they need to be productive, without the ability to do damage that hasn't been sanctioned.

Harmonic Command works across both local MCP servers running on employee machines and remote servers connected to cloud-hosted enterprise tools, giving you a single governance layer regardless of where the agent operates.

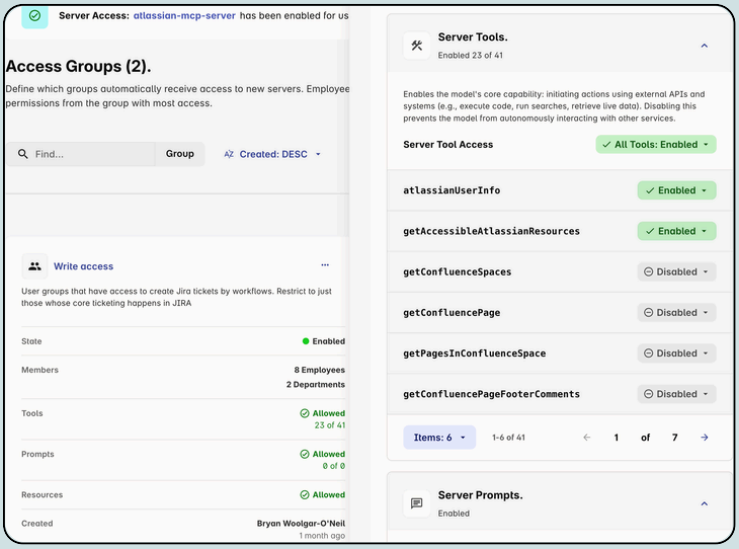

Granular Server Access Controls

Protect Sensitive Information in Agentic Workflows

The same small language models that protect human AI interactions now apply to everything agents do.

When an agent attempts to pass sensitive information to a connected tool, Harmonic detects it in real time, using the same context-aware analysis that understands intent, not just patterns.

Instead of killing the workflow, Harmonic guides the agent toward a safe alternative. The task keeps moving. Your data stays protected.

Harmonic also provides real-time protection against prompt injection attacks. When malicious instructions embedded in external content attempt to hijack an agent's behavior, Harmonic identifies the threat and blocks it before the agent can act on it, keeping your agentic workflows safe from manipulation.

Complete Logging and Vizualization of Alerts and Interventions

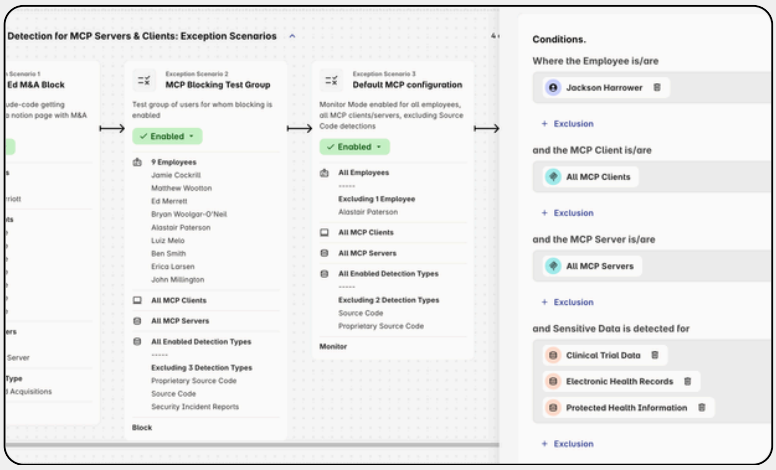

Granular Controls of What Agents Can Share

Harmonic Guide vs. Harmonic Command

Guide gives security teams real-time visibility and enforcement across browser and endpoint AI tools. Understand how your workforce uses AI, detect sensitive data in context, and coach employees toward safer behavior without slowing them down.

Command extends that same governance to autonomous AI agents. Everything in Protect, plus the ability to discover, monitor, and guide AI agents operating across your enterprise systems.

	Guide	Command
Visibility		
Shadow AI discovery (10,000+ apps)	✓	✓
Prompt-level visibility across browser, desktop and CLI	✓	✓
AI provider API call discovery and cataloguing	✓	✓
AI Usage Intelligence	✓	✓
Identity enrichment (Entra / Okta)	✓	✓
MCP client and server discovery	—	✓
Enforcement		
Real-time sensitive data detection and coaching	✓	✓
Customizable policies and groups	✓	✓
Security integrations and webhooks	✓	✓
Granular agent controls (read / write / execute)	—	✓
Prompt injection protection	—	✓

Why Harmonic Command?

Harmonic Command gives security teams governance across every AI surface, from browser to endpoint to agentic workflows on the device and in the cloud, managed from a single platform.

Intent-based action controls
Our small language models understand what users and agents are trying to do, not just the data involved. Policies are enforced based on the full context of the interaction, so you get fewer false positives and more confident decisions.

Guide agents without breaking workflows
When an agent encounters sensitive data or a risky action, Harmonic redirects it toward a safe alternative. The task keeps moving. Your data stays protected.

One platform for humans and agents
Browser, endpoint, and agentic AI governed from a single console with one policy engine. Define your rules once and they apply everywhere AI operates across your organization.

> While we did understand what people were using — because we had that visibility — we didn't know how they were using it, and the data leakage was definitely a big one. ...
>
> That's when we went looking for outside help. ... It's literally hours to seconds. If we deploy a tool and it's covered by Harmonic, we have insight right away.
>
> **Mike Janielis Principal Information Security Architect Advisor360**

Getting Started with Harmonic

Getting started with Harmonic is quick and easy. Our browser extension and lightweight endpoint agent can be installed via MDM.

The endpoint agent is used to install and configure the MCP gateway, along with providing detailed information on agent status (skills/configurations/plugins/etc.) and employee inputs.

Within 30 minutes, the platform can be rolled out to your entire organization.

harmonic

harmonic.security sales@harmonic.security

As AI expands from employee tools to autonomous agents, organizations need governance that keeps pace. Harmonic delivers AI Governance and Control (AIGC), the intelligent layer that secures and enables both human and agentic AI workflows. By understanding intent and data context in real time, Harmonic gives security leaders the confidence to let their companies move fast with AI.