MCP Gateway:

Visibility and control over agentic AI

Uncover Shadow AI Use and Gain Control

Enterprise AI has evolved from single-prompt interactions to agentic workflows. This has created massive productivity gains, but it's also introduced a new, invisible attack surface.

The Risk of MCP (and the security gap is exposes)
Agentic AI workflows are reshaping how work gets done, connecting AI models directly to company data, APIs, and systems. But this new workflow layer operates outside traditional security controls.

Sensitive data can move between AI tools and business systems without oversight, leaving security teams blind to:

- Which MCP clients and servers are in use
- What data is being exchanged
- When risky workflows occur

Without visibility and control, enterprises face data leakage, workflow hijacking, and compliance gaps. to train on customer data. No drowning in false positives.

Introducing the MCP Gateway
The Harmonic MCP Gateway is a developer-friendly, locally installed gateway that gives security teams complete visibility and control over their organization's agentic AI ecosystem.

It transparently intercepts all MCP traffic, allowing security teams to discover what clients and servers are in use, enforce granular policies to block risky actions, and—most importantly—apply Harmonic's industry-leading sensitive data models to prevent the exposure of critical intellectual property and other sensitive information.

Key Features

- **Discover agentic AI workflows,** including all MCP clients and servers across your organization.
- **Prevent sensitive data exposure in agentic workflows,** including source code, financial data, and strategic plans.
- **Enable secure AI innovation by coaching AI agents at the point of blocking,** allowing them to find safe alternative paths.
- **Enforce centralized policies** to block risky MCP servers or restrict high-risk capabilities like production database writes.

Visibility into MCP Use

Discovery & Inventory
Automatically discover and inventory all MCP clients (e.g., Cursor, Claude Code) and servers (official vendor and locally built) in use.

Usage Analytics
Understand which employees are using which clients or servers and how frequently via dashboards and saved views.

Invocation Logging
Capture detailed audit logs of every interaction for forensic analysis and compliance.

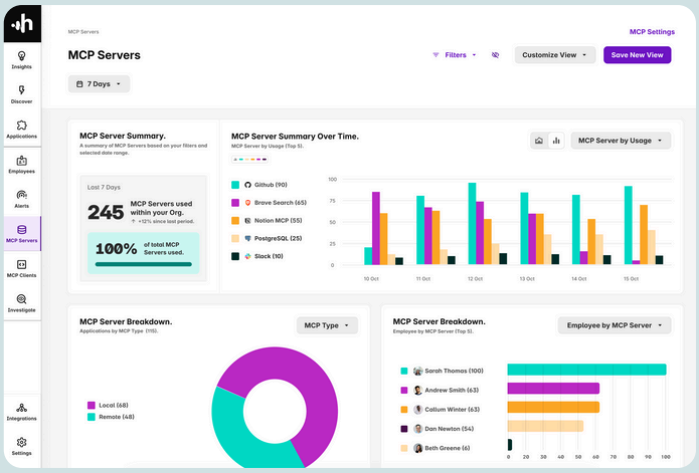

MCP Usage Visibility

Control

Centralized Policy Enforcement
Define and enforce global policies to block entire MCP servers or restrict specific high-risk capabilities (e.g. tools that can write to production databases).

Alerting & Integration
Receive real-time alerts for policy violations and sensitive data events. We integrate seamlessly with your existing security stack, ensuring alerts appear where your teams already monitor and respond (e.g. SIEM and SOAR platforms).

Sensitive Data Detection
Leverage existing Harmonic sensitive data models to inspect MCP traffic in real-time, identifying unstructured sensitive data like source code, financial projections, and strategic plans.

Intelligent Blocking & AI Coaching
When sensitive data is detected, we don't just block the action and break the workflow.

The gateway provides contextual, detection-specific feedback to the MCP client.

This coaches the AI agent on why an action was blocked, allowing it to find a safe, alternative path to complete its task, thereby reducing developer friction and enabling safe AI adoption.

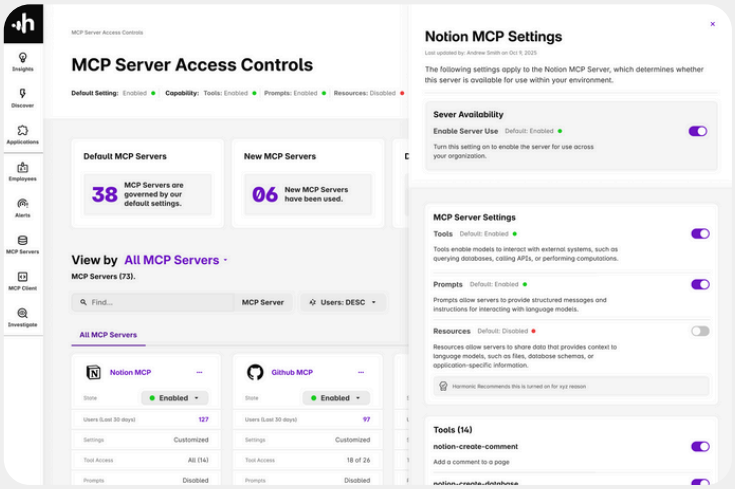

MCP Control Settings

Coach Users at the Point of Data Loss

Security teams are realizing that blocking GenAI tools is not enough; employees will use these tools regardless. By enabling end users to use these tools while nudging them against exposing sensitive data, security teams have a way to enable the business' AI use.

Harmonic's models are exceptionally fast, making accurate assessments within 200 milliseconds. This is 300 times quicker than if you were to use an LLM like ChatGPT.

When we detect sensitive data, we intercept the prompt before it is exposed and give the end user context to understand what is sensitive and how they can remove the sensitive data.

The intervention model is fully customizable in how end users may be nudged:

- Prompt user to edit, block, ignore, or redirect;
- Link to internal policies or documents;
- Define which users can upload sensitive data types to specific GenAI apps;
- Customize text, colors, and logo.

This helps to engage with users where they are at, minimizing friction and encouraging secure adoption of GenAI.

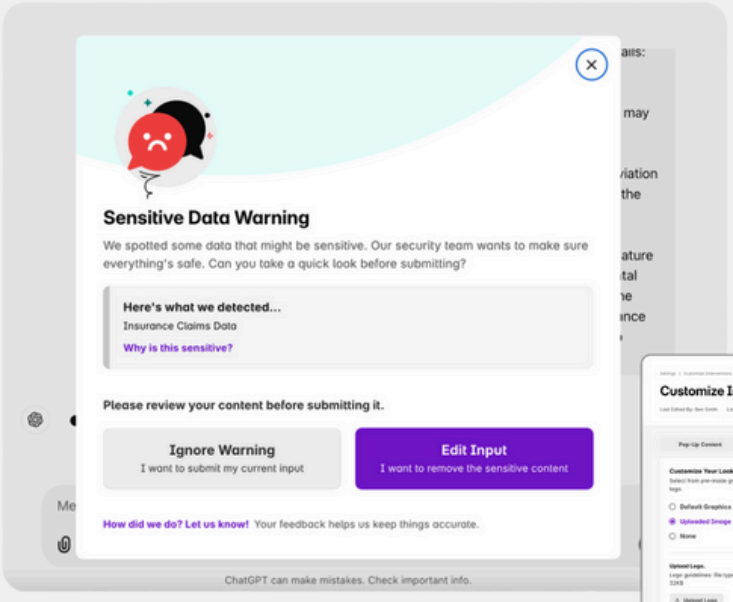

Customizable Model Displayed to End Users

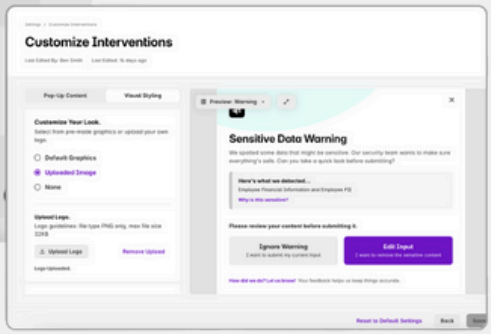

Customize Interventions and Nudges

Why Harmonic?

With the Harmonic MCP Gateway, you get instant visibility and protection without the complexity of a high-friction agent rollout.

Installation takes minutes and delivers value on day one. No heavy infrastructure or configuration required. The lightweight gateway runs seamlessly on Windows, macOS, and Linux, giving every team the same secure foundation for agentic workflows.

Securely Accelerate AI Innovation
The rise of agentic AI is inevitable, but it doesn't have to be risky.

With Harmonic MCP Gateway, enterprises can innovate with confidence, maintaining full visibility and control over how AI interacts with their systems and data.

"Every security leader I know is trying to get ahead of AI-driven workflows.

It's exciting to see Harmonic tackling this head-on, so teams can be confident to innovate safely."

Mike Janielis Principal Information Security Architect Advisor360

Getting Started with Harmonic

Getting started with Harmonic is quick and easy. Simply install the Harmonic MCP Gateway to start gaining insights into Agentic AI workflows and secure your sensitive data.

Within 30 minutes, the Harmonic MCP Gateway may be rolled out to your entire organization with Group Policy Object (GPO), Microsoft Intune, JAMF or Kandji.

Get Started >

harmonic harmonic.security sales@harmonic.security

As every employee adopts AI in their work, organizations need control and visibility. Harmonic delivers AI Governance and Control (AIGC), the intelligent control layer that secures and enables the AI-First workforce. By understanding user intent and data context in real time, Harmonic gives security leaders all they need to help their companies innovate at pace.