harmonic | Harmonic Guide

Harmonic Guide: Enable AI Use From the Browser to the Desktop

Safe Employee AI Usage Everywhere

Security teams are under pressure to enable AI and protect sensitive data at the same time. With AI now embedded in browsers, desktop apps, and developer tools, governance needs to cover every surface.

Harmonic Guide at a Glance

Harmonic Guide gives security teams real-time visibility and enforcement across browser-based AI tools and endpoint applications. 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 coaches the employee toward safer behavior, all within milliseconds.

Built for the Age of AI Agents

AI exists everywhere employees work. Developers use Claude Code and GitHub Copilot. Marketing teams create assets in Canva. Business analysts run data through ChatGPT Desktop. Lawyers summarize contracts in Microsoft 365 Copilot.

Harmonic Guide governs all of it with one unified policy engine, one set of detections, and one console. No blind spots between surfaces, no gaps between tools.

Key Features

- **Prompt level visibility across 1,000+ apps**, including tools employees adopt without IT approval.
- **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 with inline interventions** that warn, nudge, or block at the point of data loss.

Why People Love Harmonic

Understand intent, not just patterns. Our models detect sensitive information in context, no reg(ex) needed.

Gain ACTIONABLE insight into how AI is used. See where ROI is being delivered and where risk exists.

Operate at employee speed. Inputs are analyzed and interventions are delivered in milliseconds.

Guide employees instead of blocking them Real-time coaching explains what was flagged and why, building security awareness with every interaction.

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.

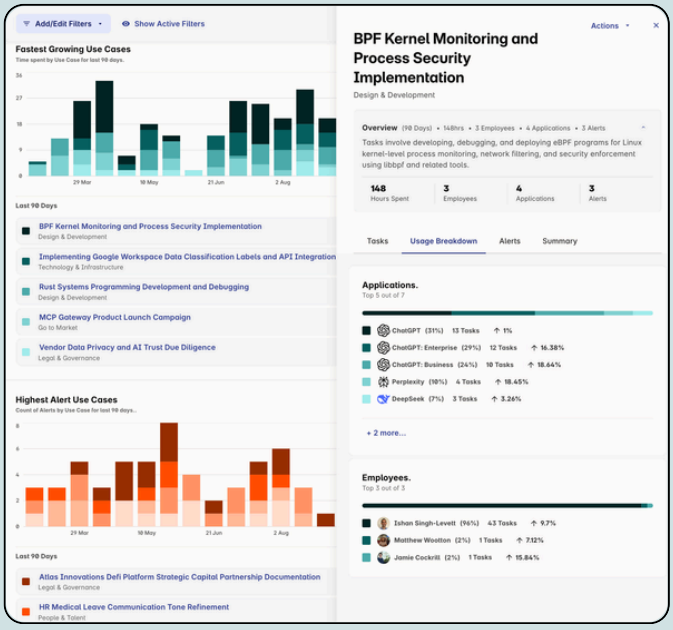

Actionable Usage Insights

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.

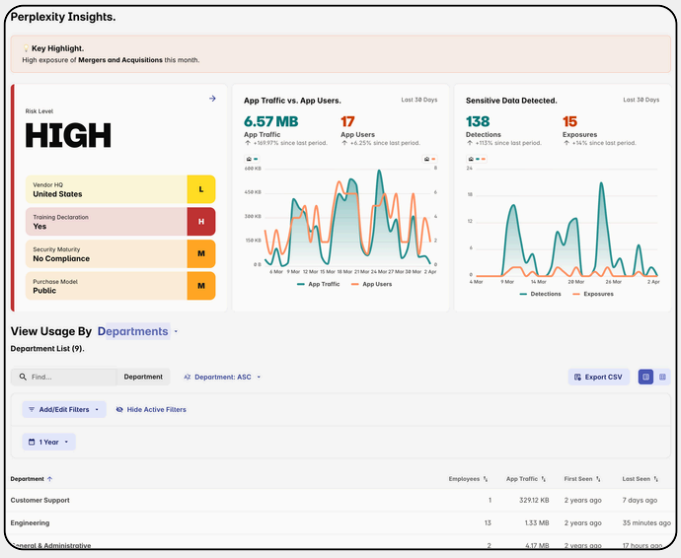

Complete Visibility by Application

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.

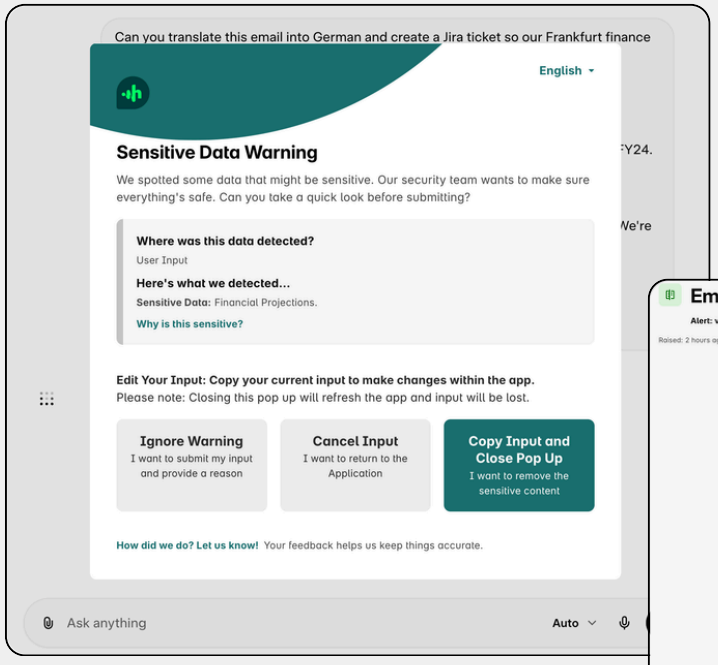

Customizable Notice Displayed to End Users

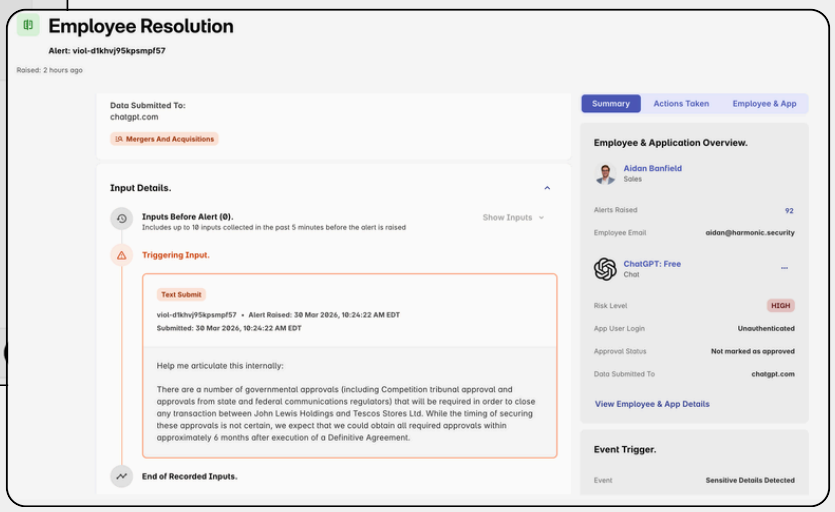

Complete Data of Each Alert and Resolution

Why Harmonic Guide?

Harmonic Guide empowers security teams to safeguard sensitive data whererever employees leverage AI.

Deep understanding of AI use cases
Usage Intelligence maps AI activity to real business use cases, giving security, IT, and AI leaders the insight to make informed decisions.

Inline models that coach in real time
Small language models read the full interaction, not just keywords or patterns, and deliver taliored interventions in under 200 milliseconds.

A platform that grows with you
Start with browser and endpoint governance today. When AI agents and MCP workflows enter the picture, Harmonic scales with you.

> 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. Simply install the browser extension or lightweight endpoint agent to get started.

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

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.