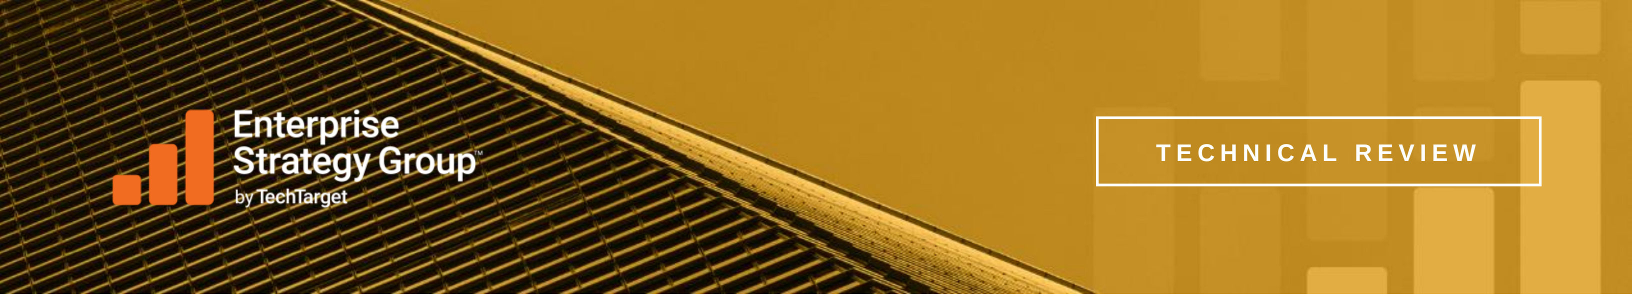

Enterprise Strategy Group by TechTarget

TECHNICAL REVIEW

NOVEMBER 2024

Take Advantage of Generative AI Securely With Harmonic Protect

Justin Boyer, Technical Validation Analyst

Abstract

This Technical Review by TechTarget's Enterprise Strategy Group details our validation of Harmonic Security's Harmonic Protect platform. We evaluated how Harmonic Protect helps organizations safeguard sensitive data and prevent its spread and misuse, especially via generative AI applications.

The Challenges

The development of generative AI has fundamentally changed how businesses operate and how employees work. While generative AI provides new tools to accelerate business operations, it also introduces new risks, including the risk of data leakage when employees expose sensitive data via generative AI prompts.

Organizations have already begun developing governance structures for generative AI. According to Enterprise Strategy Group research, 96% of survey respondents are either already enforcing a specific governance structure for the use of generative AI or are currently developing one.[1]

The research also shows that organizations are working hard to catch up with their employees' use of generative AI, with 86% reporting that they have or are blocking access to one or more generative AI sites. Further, 82% are concerned about data leakage as employees increasingly use generative AI, and another 82% are unsure if any employees are accessing generative AI sites or what they are doing on the sites (see Figure 1).

[1] Source: Enterprise Strategy Group Research Report, *Generative AI for Cybersecurity: An Optimistic but Uncertain Future*, April 2024. All Enterprise Strategy Group research references and charts in this Technical Review are from this report.

Figure 1. Organizations Are Cautious With Generative AI

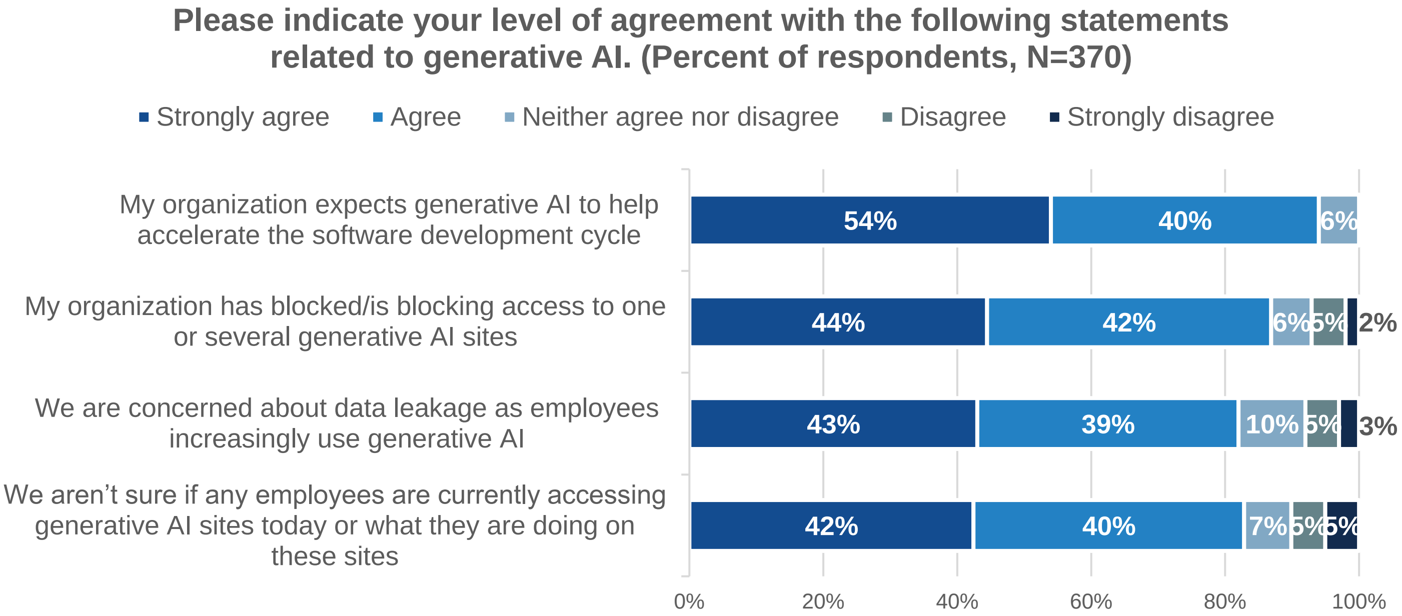

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Legacy data protection techniques, such as data labeling, creating complex data protection rules, and blocking access to services employees might use, aren't adaptive enough to the changing AI and data landscape. Organizations need a solution that protects sensitive data without completely shutting down employees or ignoring the benefits of generative AI applications.

The Harmonic Protect Platform

Harmonic Protect is a data protection platform designed to prevent sensitive data leakage, without the need for extensive data labeling or rule creation, whether that data is structured, unstructured, or semistructured. The turnkey solution features prebuilt data protection models that identify sensitive data without regular expressions. These models can assess the context of data usage, leading to humanlike decisions about whether the data is sensitive and whether its exfiltration could create undue risks to the business.

With a simple browser extension, Harmonic Protect works where the end user works (see Figure 2). Using these technologies, Harmonic can coach end users by detecting potential data loss within milliseconds and prompting the user to reconsider what data they're sending to generative AI applications. Harmonic integrates with an organization's identity provider and uses web hooks for additional functionality, such as sending messages to the employee via collaboration software or redirecting them to security awareness training platforms like KnowBe4 if they repeatedly ignore data protection warnings.

Figure 2. Harmonic Protect Platform

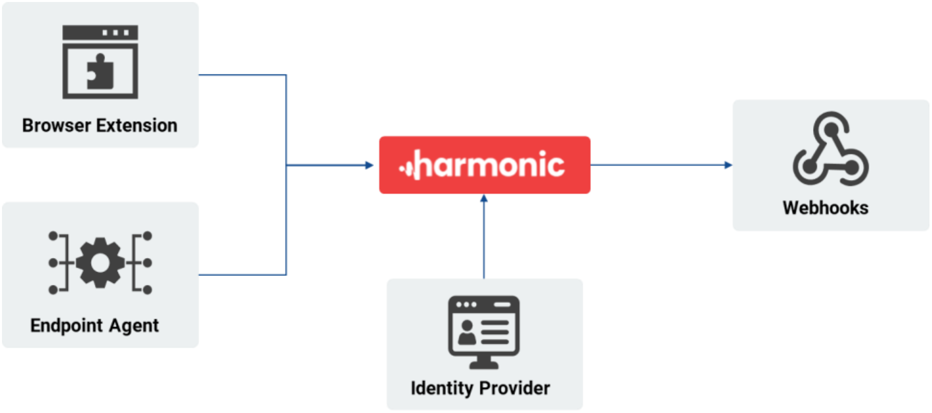

Source: Harmonic and Enterprise Strategy Group, a division of TechTarget, Inc.

Enterprise Strategy Group Tested

Harmonic Security significantly saves time for security teams, compliance teams, and AI committees. Typical data protection strategies are time-consuming and expensive, requiring extensive data labeling and rule creation, with regular expressions to try to locate sensitive data. This often leads to false positives and creates more administrative overhead.

Harmonic Protect features ready-made data models built to find sensitive data such as employee personally identifiable information (PII), employee financial information, source code, and insurance claims data. This structured, unstructured, and semistructured data can be difficult to contain. Harmonic's prebuilt models have been trained to use context to determine the risk of a particular set of data being exposed via generative AI tools such as ChatGPT, Microsoft Copilot, and Google Gemini. The prebuilt models can be activated on an as-needed basis by administrators based on business requirements.

Figure 3 shows the Employee Insights Dashboard, where administrators can view high-level statistics such as application activity, potential sensitive data breaches, and any actions taken to intervene.

Figure 3. Employee Insights Dashboard

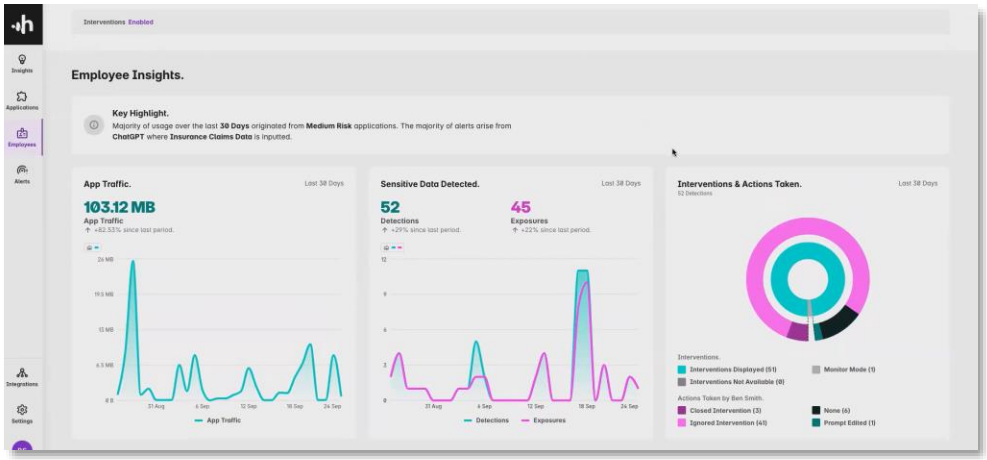

Interventions Enabled

Employee Insights.

Key Highlight.
Majority of usage over the last **30 Days** originated from **Medium Risk** applications. The majority of alerts arise from **ChatGPT** where **Insurance Claims Data** is inputted.

App Traffic. Last 30 Days
103.12 MB App Traffic
↑ +82.53% since last period.

Sensitive Data Detected. Last 30 Days
52 Detections
↑ +29% since last period.
45 Exposures
↑ +22% since last period.

Interventions & Actions Taken. Last 30 Days
52 Detections

Interventions.
- Interventions Displayed (51)
- Monitor Mode (1)
- Interventions Not Available (0)

Actions Taken by Ben Smith.
- Closed Intervention (3)
- None (6)
- Ignored Intervention (41)
- Prompt Edited (1)

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Harmonic Protect also saves time by automatically cataloging applications. Every time an employee uses an application, Harmonic will make the application's profile available within the environment. Application profiles provide details on what the application is, how it's used, and potential risks. By eliminating "shadow AI" applications in use by employees without the organization's knowledge, Harmonic reduces risk and provides transparency around how these generative AI applications use the data within their prompts. Many companies are wary of generative AI applications that use submitted data to train their models, but because Harmonic Protect provides application profile details, administrators can decide which applications to allow and which to control more closely. Additionally, this reduces the load on compliance teams that are increasingly tasked with creating AI inventories for the organization.

Harmonic Protect detects potential sensitive data exposure using its prebuilt models. Administrators can turn the models on and off based on business requirements and create custom workflows to define actions based on the scenario. Figure 4 shows the Detection Catalog, where administrators can view and configure the various models. Each model entry lists the data definition, which outlines the criteria Harmonic uses to determine if a prompt contains that data type. This transparency gives organizations visibility into the models, unlike many AI-based solutions, and can also reduce false positives and time required to investigate alerts.

Figure 4. Harmonic Protect Detection Catalog

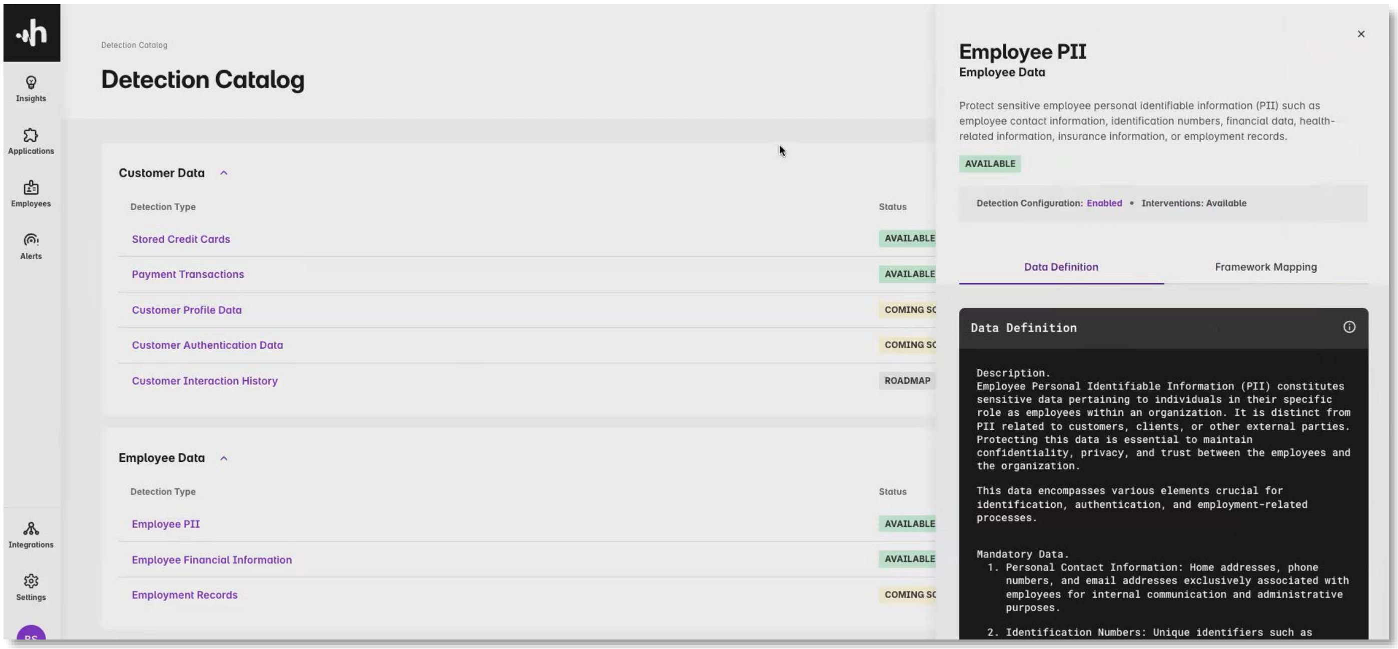

Detection Catalog

Detection Catalog

Customer Data

Detection Type	Status
Stored Credit Cards	AVAILABLE
Payment Transactions	AVAILABLE
Customer Profile Data	COMING SC
Customer Authentication Data	COMING SC
Customer Interaction History	ROADMAP

Employee Data

Detection Type	Status
Employee PII	AVAILABLE
Employee Financial Information	AVAILABLE
Employment Records	COMING SC

Employee PII

Employee Data

Protect sensitive employee personal identifiable information (PII) such as employee contact information, identification numbers, financial data, health-related information, insurance information, or employment records.

AVAILABLE

Detection Configuration: Enabled • Interventions: Available

Data Definition | Framework Mapping

Data Definition

Description.
Employee Personal Identifiable Information (PII) constitutes sensitive data pertaining to individuals in their specific role as employees within an organization. It is distinct from PII related to customers, clients, or other external parties. Protecting this data is essential to maintain confidentiality, privacy, and trust between the employees and the organization.

This data encompasses various elements crucial for identification, authentication, and employment-related processes.

Mandatory Data.
1. Personal Contact Information: Home addresses, phone numbers, and email addresses exclusively associated with employees for internal communication and administrative purposes.
2. Identification Numbers: Unique identifiers such as

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Coaching Users by Helping Them at the Point of Data Loss

Harmonic detects sensitive data within milliseconds and intercepts it before it leaves the business. Enterprise Strategy Group observed how Harmonic successfully stopped an employee from entering an insurance claim into a generative AI prompt. As shown in Figure 5, Harmonic prompted the employee to update the prompt to remove the sensitive information and included a detailed explanation of why Harmonic flagged the prompt. Harmonic flagged this data despite the absence of a concrete indication, such as a policy number. Using regular expressions to find policy numbers has limitations; without Harmonic's AI models backing them up, this sensitive data could have been leaked.

Figure 5. Employee Prompted to Remove Sensitive Information From a Prompt

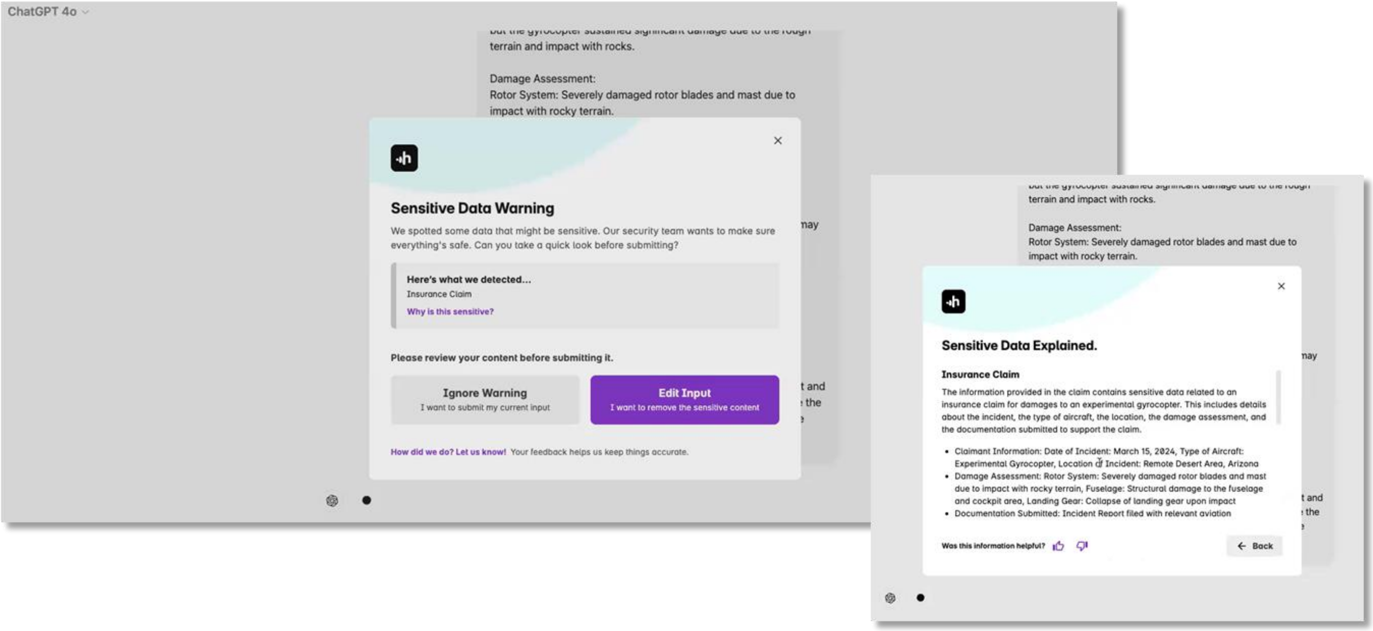
ChatGPT 4o
×
Ignore Warning
I want to submit my current input
Edit Input
I want to remove the sensitive content
×
Was this information helpful?
Back

terrain and impact with rocks.

Damage Assessment:
Rotor System: Severely damaged rotor blades and mast due to impact with rocky terrain.

Sensitive Data Warning

We spotted some data that might be sensitive. Our security team wants to make sure everything's safe. Can you take a quick look before submitting?

Here's what we detected...
Insurance Claim
Why is this sensitive?

Please review your content before submitting it.

How did we do? Let us know! Your feedback helps us keep things accurate.

terrain and impact with rocks.

Damage Assessment:
Rotor System: Severely damaged rotor blades and mast due to impact with rocky terrain.

Sensitive Data Explained.

Insurance Claim

The information provided in the claim contains sensitive data related to an insurance claim for damages to an experimental gyrocopter. This includes details about the incident, the type of aircraft, the location, the damage assessment, and the documentation submitted to support the claim.

- Claimant Information: Date of Incident: March 15, 2024, Type of Aircraft: Experimental Gyrocopter, Location of Incident: Remote Desert Area, Arizona
- Damage Assessment: Rotor System: Severely damaged rotor blades and mast due to impact with rocky terrain, Fuselage: Structural damage to the fuselage and cockpit area, Landing Gear: Collapse of landing gear upon impact
- Documentation Submitted: Incident Report filed with relevant aviation

Source: Enterprise Strategy Group, a division of TechTarget, Inc.

When prompted, employees have the option to update their generative AI prompt to remove the sensitive information or to ignore the warning. Harmonic records these prompts and alerts administrators when an employee ignores a warning. Administrators can also block users from uploading or choose to redirect them to approved generative AI applications or required security training if they ignore the warning.

Why This Matters

Generative AI applications have grown from a fun novelty to serious business tools. With that growth comes new security challenges. Ninety-six percent of respondents to Enterprise Strategy Group research stated that they are either enforcing or developing a governance structure around generative AI usage. Additionally, 82% stated that they are concerned about data leakage as generative AI use increases among their employees. Despite its power, employees using generative AI without safeguards could leak sensitive data via prompts to generative AI tools that train their models with user-input data. In areas governed by strict regulations, such as GDPR, such data leaks could lead to stiff penalties.

Enterprise Strategy Group validated how Harmonic Security's Harmonic Protect platform helps organizations manage the risks posed by generative AI. Harmonic Protect is easily deployed by adding a browser extension to end users' computers. Harmonic then watches and catalogs which generative AI applications employees use, providing detailed information about how the applications function and any potential risks. We saw how Harmonic intercedes directly with the end user when it detects sensitive data in a generative AI prompt by nudging the user to change what they've entered. Administrators can see which employees have ignored such warnings and redirect them to security training to help increase compliance.

Employees are going to use generative AI tools to help them accomplish their jobs more efficiently. However, they need guardrails to help protect the organization from data breaches and regulatory penalties. Harmonic Protect helps to enable end users, not by blocking generative AI applications altogether but by serving as a watchful eye that guides end users to the correct decision. By interacting directly with the employees using generative AI applications, Harmonic nudges them toward compliance while giving organizations a complete picture of how their employees employ generative AI.

A Customer Perspective

Enterprise Strategy Group also considered case studies and interviewed Harmonic customers to understand how they've benefited from using Harmonic Protect to safeguard their generative AI use.

Harmonic Protect Generates 96% Fewer Alerts Than DLP Solutions

An investment advisory company uses Harmonic Protect to prevent sensitive financial and personal data from being used in generative AI applications. Their existing solutions either blocked every generative AI application or generated too many alerts to realistically handle. The company struggled to gain actionable information from these tools, as they didn't provide a clear picture of AI model usage across the company.

> **"It would take forever with our current tooling to try and get feature parity [with Harmonic Protect] or even come close to it. We'd be writing regexes forever."**
>
> —CIO, investment advisory company

This company didn't want to ban generative AI altogether, so they decided to deploy the Harmonic Protect platform across their environment. After deploying Harmonic, the company is now able to pursue a "trust but verify" model. Employees can use generative AI applications to increase productivity, but Harmonic is there to prevent sensitive and confidential data from passing into the applications. The company was impressed with Harmonic's prebuilt models, which "cut through the noise" and provide a clear and concise picture of AI usage across the company. Over the previous month, Harmonic Protect had 96% fewer alerts than the existing data loss prevention (DLP) solutions, showing a massive reduction in false positives. In one case, Harmonic Protect prevented sensitive loan information from being used in a generative AI prompt. This company is convinced that no other solution would have detected this data because it didn't have any information regular expressions would detect.

Harmonic Protect Saved 75% in Implementation Costs

A hospitality services and software company based in Europe also found Harmonic Protect to be an asset in protecting sensitive data while enabling employees to use generative AI applications. This company operates in a highly regulated environment with severe penalties for data breaches. With many prominent and confidential clients, along with advanced proprietary software code and sensitive employee data, this company needs to prevent several different types of sensitive data from reaching large language models that could use it for training.

> **"Installation was easy because it was just a browser plugin. Instead of needing servers and cloud environments to begin using it, you install the browser extension and within days have it rolled out and gain visibility into the entire environment. We got up and running without a big headache, a big budget, or big consulting efforts."**
>
> —CISO, hospitality services and software company

Typically, a standard DLP solution might require expensive consultants to interview employees for weeks or months before it is designed and built. The hospitality company found that Harmonic Protect's simple and fast deployment, based on its browser extension model, saved approximately 75% on project costs, resources, and time in comparison with a standard DLP solution deployment. Another priority for this company was building a security-conscious culture. It found that Harmonic Protect's nudging of end users at the time of data loss was effective in building awareness and creating conversations around why these controls are necessary. By using Harmonic in conjunction with security awareness training, this company has seen improvements in compliance among employees and feels that a security-focused culture is within reach.

Conclusion

The acceleration of generative AI technologies has enabled increased productivity in many fields. However, its adoption has outpaced many organizations' ability to properly secure them. According to Enterprise Strategy Group research, 82% of surveyed organizations are concerned about data leakage through employees' use of generative AI applications. A further 82% indicated that they don't know which generative AI applications their employees are using or what data they might be sharing with the applications. These "shadow AI" practices lead to increased risk of data loss, data breaches, and regulatory penalties.

Enterprise Strategy Group validated that Harmonic Protect enables organizations to take advantage of generative AI while reducing risk. Harmonic uses a browser extension for web applications to monitor and report on which generative AI tools employees use and how they use them. We observed how Harmonic's pretrained AI models detected an employee entering sensitive insurance claim information into a generative AI tool and prompted the employee to remove the sensitive data. Administrators can turn these pretrained models on and off as needed as well as see what criteria the models use to determine if a generative AI prompt contains sensitive data. By intercepting the data as the user enters it, Harmonic nudges the employee to remove sensitive data from the prompt before it leaves the organization. Administrators can view employees who've ignored the warning and can send them a message via a collaboration application or redirect them to awareness training.

Instead of only reporting on what has happened and depending on complex regular expressions for detection, Harmonic Protect helps employees make better decisions in real time, with fewer false positives. If your organization wants to take advantage of generative AI applications without risking sensitive and confidential information leakage, Enterprise Strategy Group recommends you consider Harmonic Protect.

About Enterprise Strategy Group
TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

contact@esg-global.com
www.esg-global.com