harmonic

# Secure GenAI Adoption

## A CISO Guide to Overcoming Data Privacy Challenges

2024

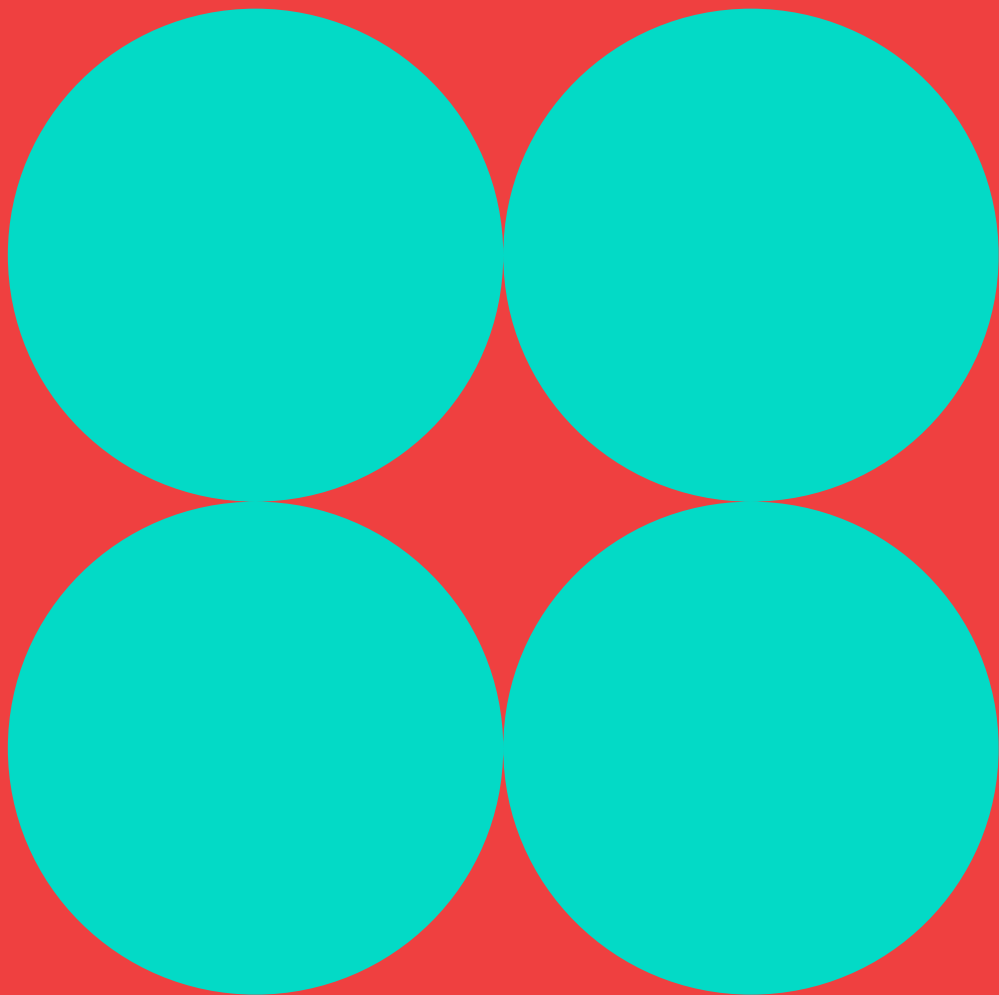

# Table of Contents


Executive Summary 2

The Value of GenAI to the Enterprise 3

Data Privacy Concerns Hold Back GenAI Adoption 4

Inventory of Applications and Use Cases 5

Policy and Guidance 6

User Training and Awareness 7

Vendor Selection 8

Process Deviations 9

NIST CSF 10

# Executive Summary

**Generative AI (GenAI) promises colossal productivity gains, but its adoption is being held back by data privacy and security concerns. In this paper, we offer practical guidance for security teams to overcome these barriers to adoption.**

"Data privacy and security" encompasses several issues: stopping intellectual property leaks into AI models, controlling the surge in apps and the danger of unauthorized AI, adhering to existing data privacy laws, and preventing employees' excessive access to sensitive information.

Addressing these security challenges demands a blend of process, people, and technology. Given the rapid emergence of GenAI, we need solutions quickly. Temporary blocks on websites are insufficient; securely embracing GenAI is vital for enterprise success.

However, growing board-level interest presents an opportunity. If security leaders can overcome these hurdles, they can position themselves as innovation enablers.

In this guide, we're sharing the best practices we've seen from speaking with hundreds of CISOs tackling these issues today. You can expect to gain practical advice on AI usage policies, AI working groups, end-user training, vendor assessment, and sensitive data detection.

You'll notice that there's very little on securing your own AI models – that's intentional. There are plenty of emerging frameworks and guides for securing your own LLMs. This guide is focused on how businesses can securely adopt GenAI.

> "AI adoption has moved from a threat to a must-have in the Board-CISO conversation."
>
> **Jerry Perullo, CISO**

### Key Contributions

This guide is a culmination of hundreds of conversations with security leaders across the world. Special thanks, however, is extended to the following individuals:

- Aaron Schaub, CISO
- Chris Douglas Vice President, Information Security, Corporate Security, and Fraud
- Jerry Perullo, Former NYSE CISO
- Mark Sutton, CISO
- Rinki Sethi, CISO, Bill.com
- Ross Mckercher, CISO
- Sascha Maier, CISO at SV Group

# The Value of GenAI to the Enterprise

## AI Washing or Productivity Boon?

It's easy to feel that AI marketing has taken over. Everywhere we turn we are force-fed another dollop of GenAI. For the most part, every existing vendor is rebadging their existing technology 'AI' or rushing out a simple chatbot on top of their existing offering.

Despite this "AI washing", early indications suggest that the productivity gains associated with GenAI are real. McKinsey anticipates that generative AI could increase labor productivity by 0.1% to 0.6%.[1] This equates to a staggering $4.4 trillion added to the global GDP every year.

GenAI adoption is expected to soar over the next 12 months. According to Gartner, 55% of organizations have already implemented or are piloting GenAI. Security teams are under pressure to adopt GenAI.[2]

## Understanding GenAI Use Cases

To truly understand the promise of GenAI, however, we need to peel away the marketing layers and understand the business use cases it can support. We analyzed the top 1,000 GenAI sites to shine a light on some of the early most popular use cases. The top five Harmonic discovered were:

- Customer Service
- Copywriting
- General Productivity
- Coding Assistant
- Email Assistant

These align with what we see as the most popular GenAI use cases right now, alongside vertical-specific use cases.

**55%**

Amount GenAI will add to global GDP annually.
**McKinsey[1]**

**$4.4T**

Of organizations have implemented or are piloting GenAI.
**Gartner[2]**

## Adoption of Commercial Offerings

To date, much of the focus around GenAI security focuses on the risks of building and deploying your own GenAI models.

While that is highly relevant to software companies implementing GenAI in their own products and a minority of other firms, the majority of companies will choose 'off the shelf' GenAI models that they can supplement with their own data. We already see examples of organizations that have attempted to 'roll their own', only to then turn to commercial alternatives.

There is now a vast array of commercial offerings to choose from; OpenAI, Microsoft, Google, and Anthropic are the obvious ones, but there are tens of thousands of applications in use across the broader GenAI ecosystem.

Unless AI is a core competency of a company, we expect this trend to continue as Microsoft, OpenAI and others continue to develop the infrastructure around their foundational models, and specialist SaaS companies mop up the other business use cases for AI.

1. McKinsey Digital, 'The economic potential of generative AI: The next productivity frontier'
2. Gartner, 'Gartner Poll Finds 55% of Organizations are in Piloting or Production Mode with Generative AI"

# Data Privacy Concerns Hold Back GenAI Adoption

**Enterprises recognize the necessity of embracing GenAI technologies to stay competitive and reap productivity rewards, understanding that the right use cases can lead to significant gains.**

**Yet, while adoption is expected to increase over the coming years, it's currently stuttering. A recent survey from MIT and Telstra found that, despite widespread interest and experimentation with GenAI, only 9% of respondents said they had adopted the technology widely.[3]**

## Barriers to Adoption

In the same MIT survey, 77% of the respondents cited regulation, compliance, and data privacy as key barriers to the rapid employment of generative AI.[4] Gartner's research tells a similar story, with "uneasiness over data security, privacy, or compliance" listed as the top barrier to adopting GenAI.[5]

Employees are piling data into these tools, and security teams have no idea what data is sensitive. Any leak of intellectual property can prove costly, yet existing tools are not built to detect sensitive data shared within vast flows of unstructured text.

On top of this, regulations are causing a headache. While many of the new AI regulations only apply to those building with AI, there are existing data privacy regulations, like GDPR and CCPA, that will have to be revisited.

## Between a Hard Place and a Block

Faced with such challenges, the knee-jerk reaction from many enterprises has been to block GenAI tools. We've already seen this play out with the likes of Samsung, Amazon, and the New York Times reportedly blocking access to GenAI tools entirely.

The appeal of these tools is so strong that employees are bypassing restrictions to use them. This leads to an unintended consequence where users resort to personal devices and networks, moving away from secure corporate systems. This has led to an extension of pre-existing Shadow IT challenges.

## Emerging Best Practices

We need to quickly establish best practices for this AI-first era. In the following pages, we will showcase top recommendations and practices from security leaders.

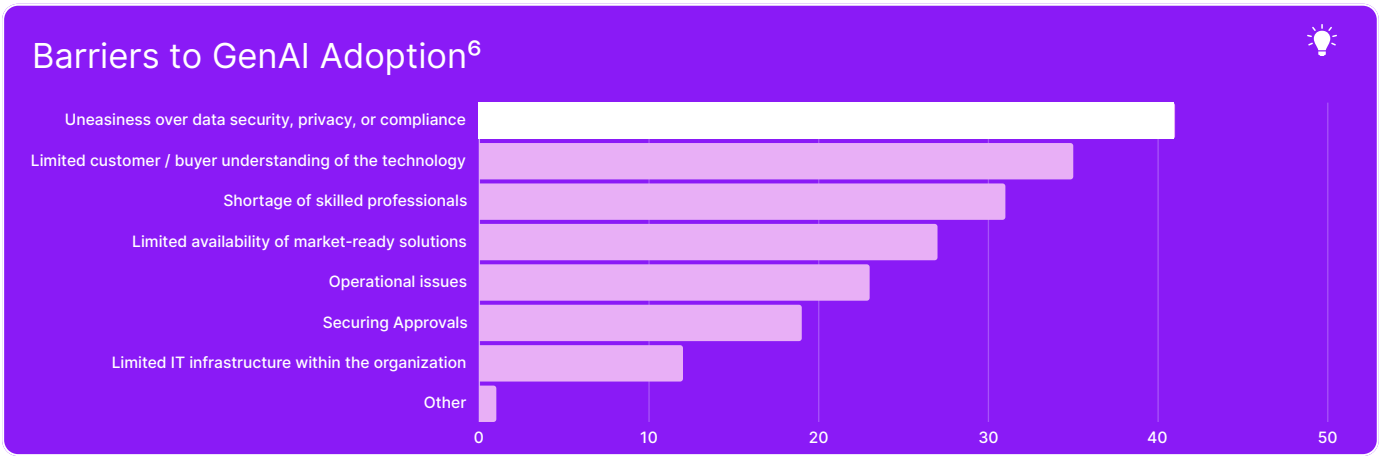


3. MIT Technology Review Insights (MITTR), 'Generative AI: Differentiating disruptors from the disrupted'
4. Ibid
5. Gartner, 'Crossing the Chasm: Tech Provider Plans for Generative AI in 2024'
6. Ibid

# Inventory of Applications and Use Cases

**The first step is to understand what people are using and what they are using it for. You cannot write a good AI policy or train end users effectively without doing so.**

**This initially mapping applies to both GenAI applications and use cases.**

## Application Audits

Application audits are essential for verifying that an organization's software is authorized, secure, and performing optimally. These audits should create a GenAI asset overview that will form a critical part of your GenAI program. Furthermore, with a surge in Shadow IT from GenAI tools, you may discover you're using far more than you think.

As well as an initial effort to map GenAI usage, organizations should aim to conduct these audits at least twice a year to stay on top of the technology being used. There are several ways to get visibility of this:

- **Speak to People!** Get to know teams and what they are using. Raise awareness and help them understand why something is a risk.
- **Manual.** Use traditional proxy/internet gateway or NGW (Next generation firewall) tools to get report data and manually filter.
- **Identity Provider.** If using Google or similar, look to see what apps "Sign-in with Google" has been used on.
- **Specialized Tooling.** Leverage third-party tools to detect Shadow IT or Shadow AI (including, but not limited to, Harmonic).

If we are to avoid making the same mistakes as with Shadow IT, we need to be more receptive to user needs. Having a tech radar that follows a "adopt, trial, assess, hold" methodology can work well.

## Use Case Mapping

Understanding use cases might be one of the most important of any GenAI program. Companies taking GenAI seriously have already begun to map out their use cases to the different assets discovered. JP Morgan, for example, already has 300 AI use cases in production. It will be the role of the AI working group to define and capture these use cases.[7] Realistically, many organizations we speak to are spending 2024 experimenting with different use cases to see what works, before wider deployment in 2025.

The exact use cases will vary from organization to organization. For example, insurance companies are using GenAI to assist with processing insurance claims. Insurance is not known for readily adopting new technologies, but these types of use cases offer huge potential for the business potential.

Several use cases stand out regardless of the organization or industry. Automated customer support, marketing content creation, meeting notes, coding assistants, and email productivity tools are extremely popular.

There's a lot to consider when mapping AI use cases. The security team must understand the use cases the business is looking to implement so they can seek a secure compliant way to implement the associated solutions. For example, by understanding these use cases, organizations can proactively recommend solutions to end users that best solve whatever business problems they're looking to solve. This will help end users achieve their goals without forcing them to use unapproved applications.


7. Joe Lin, 'AI Transformation for Financial Services (JPMorgan Chase) via Logical AI Collaborative Innovation'

# AI Usage Policies and Guidance

## Creating an AI Usage Policy

Not as many companies have AI usage policies as one might expect. According to a Gartner survey, only 33% of IT leaders reported having such guidelines. From the employees' perspective,[8] the situation appears bleaker, with just 26% believing their companies have an AI policy.[9] While two-thirds of organizations might be delaying the official launch of their AI policies, taking a proactive approach is advisable.

Policies are often interpreted as very high-level documents written by the legal and compliance team. However, to successfully overcome today's challenges, AI policies need the right stakeholders to come together, align with business objectives, and create real guidance tailored to specific use cases.

> An AI policy has value even if your organization has no immediate plans to use AI. AI will seep into our organization despite formal initiatives to embrace it. Building awareness and rules of engagement for AI might prevent an awkward situation later."
>
> **Aaron Schaub, CISO**

## Beyond the Block

Simply blocking GenAI apps and telling users to avoid uploading sensitive data to "allowed" ones is not good enough. With overzealous policies and blanket bans on AI tools, we risk forcing users underground to use unknown tools with unknown consequences. Similarly, just creating a legal document helps no one outside of the legal and compliance team.

## Including the Right Stakeholders

Identifying and committing the right stakeholders (and their time) is critical to create these policies. AI policies require more than the security or legal teams.

AI working groups are an increasingly popular way to bring these individuals together. These are cross-functional groups, typically led by a standout individual who acts as the chair. This structure allows for shared responsibility and collective decision-making, with the possibility of rotating the leadership role among committee members. Organizations that develop their own AI solutions may opt to have a Head of AI to run this group and program, but this is unnecessary for the majority of organizations.

Most areas of the business can be powerful stakeholders; security, legal, engineering, marketing, corporate communications and identity teams all have a lot to contribute.

## Building on Use Cases for Ongoing Guidance

Policies are often viewed as a static documents that employees read once for compliance reasons. However, to be effective, there must be an element of guidance that helps users adopt GenAI tools while understanding the risks. Policies should work for users, not against them.

To achieve this, AI policies must refer to specific use cases for different functions to be effective. For example, if a marketing manager knows exactly what types of content they can share with GenAI tools, this will enable specific and helpful guidance. In this regard, policy creation and employee training ought to go hand-in-hand.

8. 2023 Gartner IT Leader Poll on Generative AI for Software Engineering.
9. The Conference Board, "Survey: Majority of US Workers Are Already Using Generative AI Tools--But Company Policies Trail Behind"

# User Training

**Once you have created your AI Usage Policy, it's important to effectively train and guide end users. This is critical to ensuring that the policy is not left to collect dust on the shelf, and is actually used by end users.**

**It's important to cater your training to different maturities, departments, and use cases, helping users contextualize the guidance.**

## What to Include

As a baseline, end users should be given context on what GenAI is and how it can help or hurt the organization. However, training for GenAI security should go beyond the standard security awareness training. Here are five areas to consider:

**1) Specific approved applications.** Provide a clear list of approved GenAI tools they may use. If they visit a website that is blocked, make sure they know the approved alternative or know the process for requesting new tools.

**2) Guide them on use cases.** Show the end user examples of successful usage of GenAI applied to their role.

**3) Advise on best practices.** Show the end user examples of successful usage of GenAI. Don't begrudgingly let them use a tool; empower them to master it.

**4) Use securely.** Provide a clear list of acceptable and unacceptable uses of GenAI tools, with a focus on what data can or cannot be used in their specific role.

**5) Clarification.** Provide mechanisms for real-time clarification or Q&A when they face ambiguous or evolving GenAI use situations. This could be as simple as Slack or Teams channel.

## Delivery Methods

Requiring users to click through a questionnaire or attend a 1-hour videoconferencing call is unlikely to be engaging and effective.

Consider trying different formats that engage users. This can be achieved with gamification, personalization, and ensuring there are a variety of training delivery methods.

## Embracing a New Approach

Amidst all the security challenges of GenAI, there's a huge opportunity to shift our approach to end users. Instead of detecting policy violations that "catch" users, instead focus on coaching those users to better use GenAI securely.

When you notice an employee using a new tool or sharing something sensitive, use this as an opportunity to a) understand their needs b) train them on security risks c) offer an alternative.

Based on these learnings, build a process that will support the organization in adapting new Gen AI services in a structured way.

# Vendor Selection

**While employee use cases and needs are a vital driver for bringing on a new vendor, it's important to remember that this will ultimately be a enterprise decision.**

**Any effective GenAI security program must ensure that there are processes in place for creating new suppliers and ensuring those suppliers are handling your data securely and are not using your data to train their models.**

### Data Retention and Training Policies

All the traditional rules of supply chain security apply to GenAI, but there are additional areas of consideration that require special attention.

Data training policies, in particular, require scrutiny. This can be challenging as these policies are often opaque, mixed in with other policies, or hidden. Vendors rarely want you to dig into these details, but it's important to do so.

To add to this complexity, existing vendors that have already been approved may change their terms and conditions around training on your data, sometimes with very limited communication or visibility.

This should be a key consideration when selecting GenAI providers, and so we have some suggested questions to ask below. These are not an exhaustivelist for third party risk, and should be used in conjunction with traditional assessments.

| Key AI Supply Chain Management Considerations | |
|---|---|
| What foundational model do you use? | Understand whether your data is being sent to OpenAI, Anthropic, or other third parties with associated implications. Note that if the vendor is using open-source models such as Llama and Mixtral you will want to understand more about the security measures they have in place around the model (see next question). |
| What do you have in place to protect against OWASP's Top Ten Risks? | OWASP's Top Ten is a good starting point for understanding how the vendor is protecting against GenAI security issues that could pose a risk. |
| Describe the data that your model was trained on including an explanation of the use of any copyrighted material | This will give you an idea of both the uniqueness of the vendor's offering, but more importantly, whether there could be copyright issues in using the outputs of a model trained on dubious data sets. For more on this, see the next section on SBOMs. |
| What are the data retention policies? | Self-explanatory |
| What is your policy regarding training your GenAI model on our company data? | It is critical to understand whether your sensitive data might end up training someone else's model, with the potential for it to be discoverable by other users of that model. |
| Does it adhere to established security compliance frameworks? | With emerging regulations, such as the EU AI Act, ensure the vendor is adhering to these . |

# Process Deviations

**It's one thing to have an AI policy and another to enforce it. In this section, take two of the most popular components of AI policies: use approved GenAI applications, and do not upload sensitive data.**

## Detect Use of Unapproved Applications

While some companies are fully leaning forward and allowing any GenAI apps, most companies have blocked some number of them. Blocking can be a sensible approach – especially if this decision is based on the tool training on your company data.

Either way, once you have defined this in your GenAI Usage Policy, you'll likely want to detect employees in violation of the policy. This can be done today is proxy-based allowlists and blocklists in cloud access security broker (CASB) and secure web gateway (SWG) products. These tools make it easy to block GenAI tools, albeit will a lack of granularity around specific applications and their risks.

Whichever way you choose to block applications, make sure you don't give users a dead end. Provide an alternative tool or spark up a conversation about needs and requirements.

## Sensitive Data Leakage

Some of the worst-case scenarios for GenAI risks involve exposed sensitive data or intellectual property.

Employees are piling data into tools and sensitive data leakage is inevitable. Our analysis shows that 40% of AI apps require some sort of file upload. Menlo Security tracked an 80% increase in attempted file uploads to GenAI sites.[10]

However, GenAI presents a whole new challenges. The way users interact with GenAI tools means that it's not enough to focus on protecting data within files.

> Shadow AI will spawn as organizations struggling to manage regulatory, privacy, and security issues won't be able to keep up with widespread bring-your-own-AI (BYOAI). In addition to just genAI tools, employees will also use personally owned AI-infused software for work, adding to the BYOAI boom in the coming year."
>
> **Forrester Research**

In reality, users are creating prompts and pasting text into these tools. Traditional data protection tools are not built to detect sensitive data in unstructured data.

More effective approaches to sensitive data detection take additional context, such as behavioral insights from UEBA or identity tools or combine insights from an insider risk program.

## Taking a Fresh Approach

At Harmonic, we're putting enormous effort into providing pre-trained data protection LLMs that will detect sensitive data of any kind, in the same way a human would.

Harmonic offers a simpler alternative to complex DLP rules by allowing you to describe your sensitive data in plain English. Our models evaluate all data entering GenAI applications, considering the app, user, and data context, to effectively safeguard sensitive information.

10. Menlo Security, 'The continued impact of generative AI on security posture'

# **Securing GenAI** A Framework for Success

**To best guide security teams, we have mapped the areas in this paper, as well as others, to the NIST Cyber Security Framework 2.0.**

With the new “Govern” stage, we feel this acts as a compelling strategic handrail for security leaders to frame the security program. NIST CSF is also a great way to explain the status of the program to the Board and explain progress over time.

We’ve taken each of these six stages and provided controls and considerations at each of these for you to use as a reference guide.

Most companies we speak to are doing a reasonable job of the governance stage, but find it harder to add maturity across the other five stages.

We will soon be publishing a handbook with more detailed guidance at each of these stages

## Mapping GenAI Controls to NIST CSF 2.0

| | | | | | |
|---|---|---|---|---|---|
| **Govern** | AI Governance | AI Policy Development | Regulatory Compliance | Supply Chain Managment | Stakeholder Engagement |
| **Identify** | Use Case Mapping | Quantitative Risk Assessment | AI Intellectual Property Management | Application Audits | |
| **Protect** | Targeted Security Training for AI | AI Data Management | Secure AI Integration | Access Controls | App Blocking |
| **Detect** | Unsanctioned App Use | Sensitive Data Leakage | Vulnerability Management | GenAI Threat Intelligence | |
| **Respond** | Incident Response Plan | Crisis Management for AI Incidents | Legal and Regulatory Readiness | | |
| **Recover** | AI Data Recovery Strategies | Business Continuity Planning | Post-Incident Analysis | | |

# Thanks for reading!

**The rapid rise of generative AI (GenAI) offers unprecedented opportunities. But it also brings new threats as current data protection methods struggle to keep pace. Sensitive information can be unwittingly leaked to insecure models or applications. Left unaddressed, this opens the door to compliance violations, privacy breaches, and loss of valuable intellectual property.**

- Harmonic Security provides a unique solution:
- Total AI Visibility: Discover the full breadth of AI services in use within your organization, pinpointing hidden risks.
- Precise Risk Control: Implement targeted controls for high-risk AI applications, safeguarding critical data.
- Virtual Security Analyst: Reduce your security team's workload with automation while bolstering expertise
- Confident AI Adoption: Embrace AI's benefits knowing your data is safe.

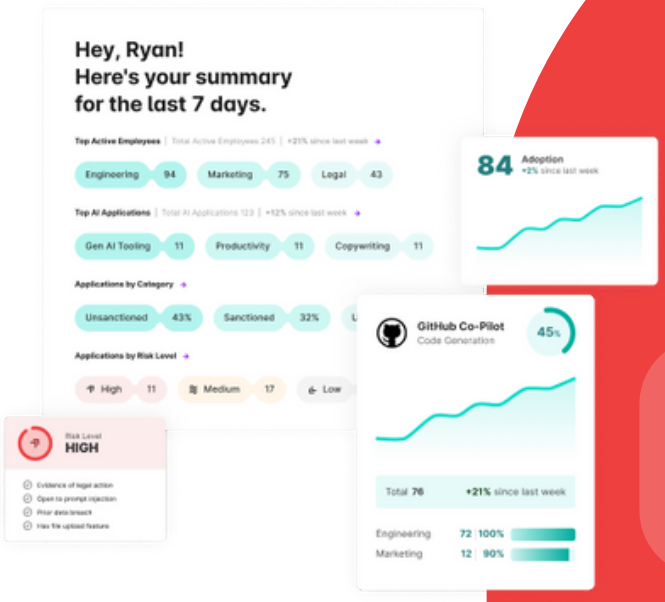


# Let's Talk!

**Request a demo**
harmonic.security


harmonic.security