USE CASE

Enterprise AI Security

AI Has Joined the Team

Enterprise AI platforms can become deeply woven into how people work. Employees use them to create content, analyze information, write code, summarize documents, and make decisions. In the process, they may share prompts, upload files, and retrieve information that contains sensitive or business-critical data.

The risk isn't necessarily the AI platform itself. It's what happens when the platform interacts with your data.

A user with legitimate access to sensitive information may still share more than they should. A seemingly routine prompt may contain confidential information. An attachment may include data that was never intended to leave its original environment. As users, permissions, data, and AI usage change, yesterday's governance decisions may no longer be sufficient.

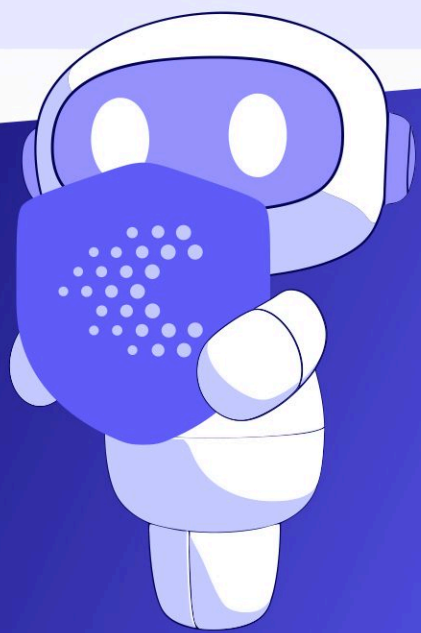
Security teams need visibility into what approved AI is doing with enterprise data—not just assurance that the application itself passed a security review.

Adoption is Just the Beginning

Sensitive data moves faster.
Employees use approved AI platforms to work with information that may be sensitive, regulated, or business-critical.

Context determines risk.
Not every AI interaction is risky. Understanding the data involved is what separates routine interactions from those that create impactful exposure.

Governance can't be one-and-done.
As data and permissions evolve and as AI usage expands, security teams need continuous visibility, not periodic assessments.



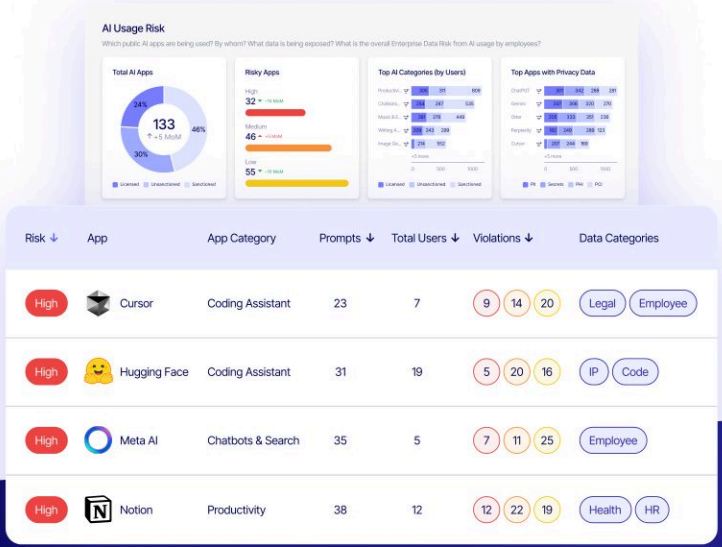
You Need Continuous Insight Into AI and Your Data

Understand how sanctioned AI interacts with sensitive information, so you can protect data in the moment, prioritize meaningful risk, and continuously adapt governance as AI becomes part of everyday work.

Accelerate Safe AI Adoption

Concentric AI provides continuous visibility into how sanctioned enterprise AI platforms interact with sensitive and business-critical data. By connecting AI activity to data, user, and organizational context, as well as sharing activity, security teams can pinpoint where meaningful exposure occurs and apply controls where they matter most. The Semantic Intelligence™ Engine understands the sensitivity, meaning, and business context of the information involved, so AI interactions can be evaluated based on the data at stake—not just the application or activity itself.

- **See what AI is doing with your data.** Track prompts, responses, and attachments across supported enterprise AI applications to understand how sensitive information is being shared.
- **Protect data before it leaves.** Apply context-aware guardrails that can warn users, block prompts, or redact confidential information before it is shared with approved AI applications.
- **Focus on meaningful risk.** Quickly identify high-risk users, departments, applications, and data flows so security teams can prioritize response efforts.
- **Keep governance current.** Continuously monitor AI activity and adapt policies as users, data, permissions, and AI adoption evolve.



Risk	App	App Category	Prompts	Total Users	Violations	Data Categories
High	Cursor	Coding Assistant	23	7	9 14 20	Legal, Employee
High	Hugging Face	Coding Assistant	31	19	5 20 16	IP, Code
High	Meta AI	Chatbots & Search	35	5	7 11 25	Employee
High	Notion	Productivity	38	12	12 22 19	Health, HR

From Deployment to Daily Operations

Continuous visibility and context-aware controls help organizations:

Protect sensitive information by identifying and preventing inappropriate exposure via sanctioned AI platforms.

Reduce investigation time by understanding the user, AI application, data, and activity involved in a potentially risky interaction.

Strengthen AI governance by creating policies informed by actual AI usage and the sensitivity of the data.

Demonstrate accountability with comprehensive records of prompts, responses, attachments, and AI activity to support investigations, governance, and compliance.

About Concentric AI

Concentric AI enables organizations to accelerate safe AI adoption through its unified AI and data security governance platform. Powered by the Semantic Intelligence Engine's patented language models, the platform discovers, governs, and protects business-critical and compliance-sensitive data at rest, in motion, and used by AI through a consistent policy framework.

Get a Demo