PLATFORM

CONCENTRIC Ai

AI Data Loss Prevention (DLP)

Your AI Prompt Has a Secret

AI makes it easy to turn a rough idea into a polished presentation, summarize a pile of documents, or solve a stubborn problem.

It can also make sharing sensitive information as easy as pasting it into a prompt. A customer record, confidential source code, or internal document shared with a public AI tool can put valuable business data at risk.

The tricky part is understanding what a data record contains. Intellectual property, confidential plans, and proprietary business details may lack recognizable patterns or keywords, making them harder for traditional DLP tools to detect.

And the AI landscape keeps expanding. Users rely on familiar tools, experiment with emerging applications, and sometimes turn to services that security teams don't even know exist. Protecting sensitive data means keeping up with what users share and where they share it.

The Prompt Is the New Leak

Not all secrets look alike.
Proprietary business information and intellectual property can slip past tools that rely on recognizable patterns.

One click can share a lot.
Prompts, copy-and-paste, and file uploads make it easy to send sensitive information to an AI application.

Shadow AI changes the game.
Unsanctioned and emerging AI tools create more opportunities for sensitive data to slip through the cracks.



Keep AI on Its Best Behavior

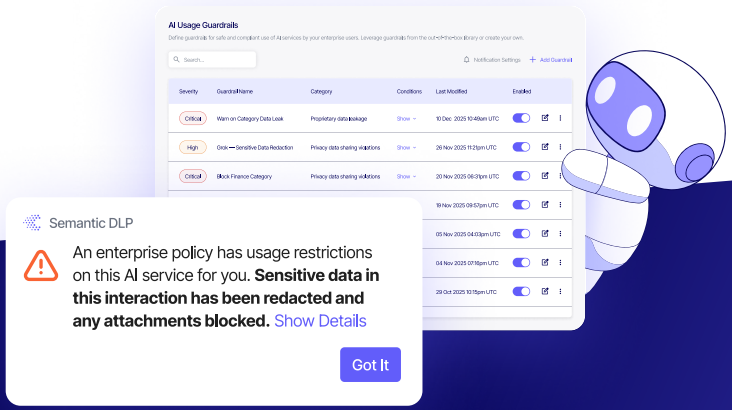
Concentric AI enforces AI policies based on what data means—not just keywords—so security teams can protect sensitive information without blocking productive AI use.

Keep the AI Flow. Stop the Data Leak.

Concentric AI's AI data loss prevention (DLP) inspects interactions in real time, identifying sensitive information before it reaches an AI application. Our Semantic Intelligence Engine uses patented language models to understand data content and context, helping identify sensitive information that pattern-based approaches may miss. Security teams can tailor guardrails by user, department, sensitive data type, and AI application. Depending on the policy, the platform can warn users, block prohibited sharing, or redact sensitive content while allowing the rest of an interaction to continue.

Protection spans browser and desktop use, not just a predefined list of popular AI tools. The browser extension applies guardrails to browser-accessed AI applications, including unsanctioned and emerging tools, while the Semantic DLP Agent extends protection to desktop apps. Concentric AI also logs prompts, responses, and policy violations across both environments to support compliance and investigations.

- **Understand what others miss.** Categorize sensitive data inline by meaning and context to identify business-critical content that pattern matching alone may overlook.
- **Set guardrails that fit the situation.** Create granular policies by user, department, sensitive data type, and AI application. Warn users, block prohibited sharing, or redact sensitive content before it reaches the AI tool.
- **Protect AI wherever it shows up.** Extend guardrails beyond familiar, sanctioned tools to help address exposure through unsanctioned and emerging applications.
- **Keep a record of AI activity.** Log prompts, responses, and policy violations to support compliance reviews and investigations.



Smarter AI. Fewer Oops Moments.

Concentric AI helps organizations safeguard sensitive data without putting the brakes on useful AI adoption.

Keep work moving by warning users or redacting sensitive content when blocking an entire interaction isn't necessary.

Reduce blind spots by helping security teams identify and address potential data exposure across a changing landscape of AI applications.

Support compliance and investigations by providing the context security teams need to understand potential data exposure.

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