CYERA

EBOOK

AI Data Protection Use Cases for Leaders

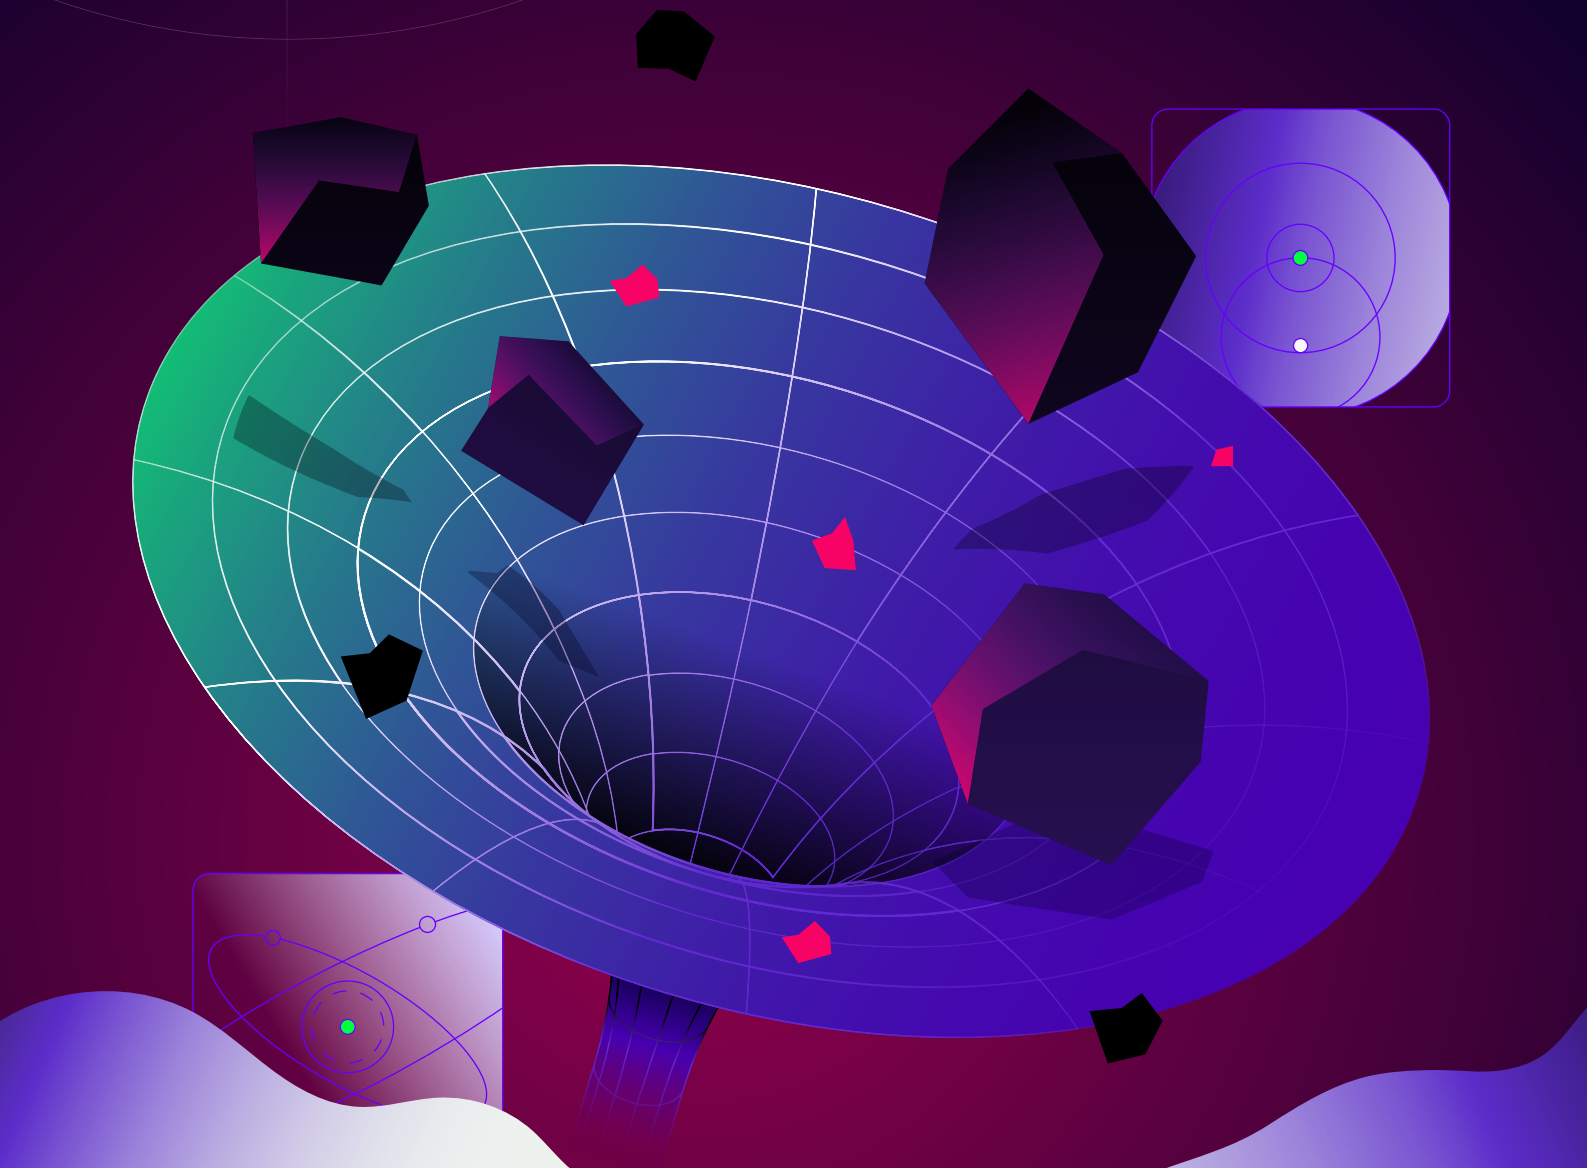

Table of Contents

The Role of CISOs and CDOs in AI Data Protection

> "80% of unauthorized AI transactions will be caused by internal violations of enterprise policies concerning information oversharing, unacceptable use or misguided AI behavior."
> - *Gartner*

AI has elevated the role of security and data leaders.

CISOs and CDOs now sit at the center of AI governance - ensuring data is protected, usage is compliant, and risk is managed at machine speed.

Boards and CEOs are turning to them to lead. And the stakes are high: 80% of unauthorized AI transactions will be caused by internal violations of enterprise policies concerning information oversharing, unacceptable use or misguided AI behavior. The lack of guardrails is accelerating how far and how fast sensitive data flows into tools like Copilot, ChatGPT, and internal models, often beyond your line of sight.

You don't need to pump the brakes. You need visibility, control, and confidence. Cyera finds and classifies sensitive data before it reaches AI tools, monitors how it's used, and flags risks in real time, so you can build boldly and securely

This guide outlines five urgent AI data protection use cases leaders are addressing now because leading the AI transformation starts with securing the data that fuels it.

Source: Gartner, Market Guide for AI Trust, Risk and Security (TRiSM), 2024

01 AI Visibility

Visibility is the foundation of everything else.

From AWS Bedrock to Azure OpenAI and Google Vertex, generative models are embedded into workflows that often operate outside of security's direct oversight. A 2024 industry survey found that 79% of IT leaders had experienced negative outcomes - such as data leakage and hallucinations - from shadow AI, while 89% of AI activity remains invisible to security tools.

Why it matters

Organizations that successfully scale domain-specific GenAI applications are more likely to outperform on ROI. But that value only materializes when data protection is built in from the very beginning.

What good looks like

- Discovery and inventory of AI-enabled tools, models, agents and applications
- Pre-training data hygiene, including auto-labeling sensitive data before it reaches AI

How Cyera Helps

AI Guardian discovers all active AI systems across the environment - whether embedded, public, or internal - and inventories the tools, users, models, and prompts involved. It tracks usage patterns and classifies the data being touched to give teams a complete view of where and how AI operates.

02 Embedded AI

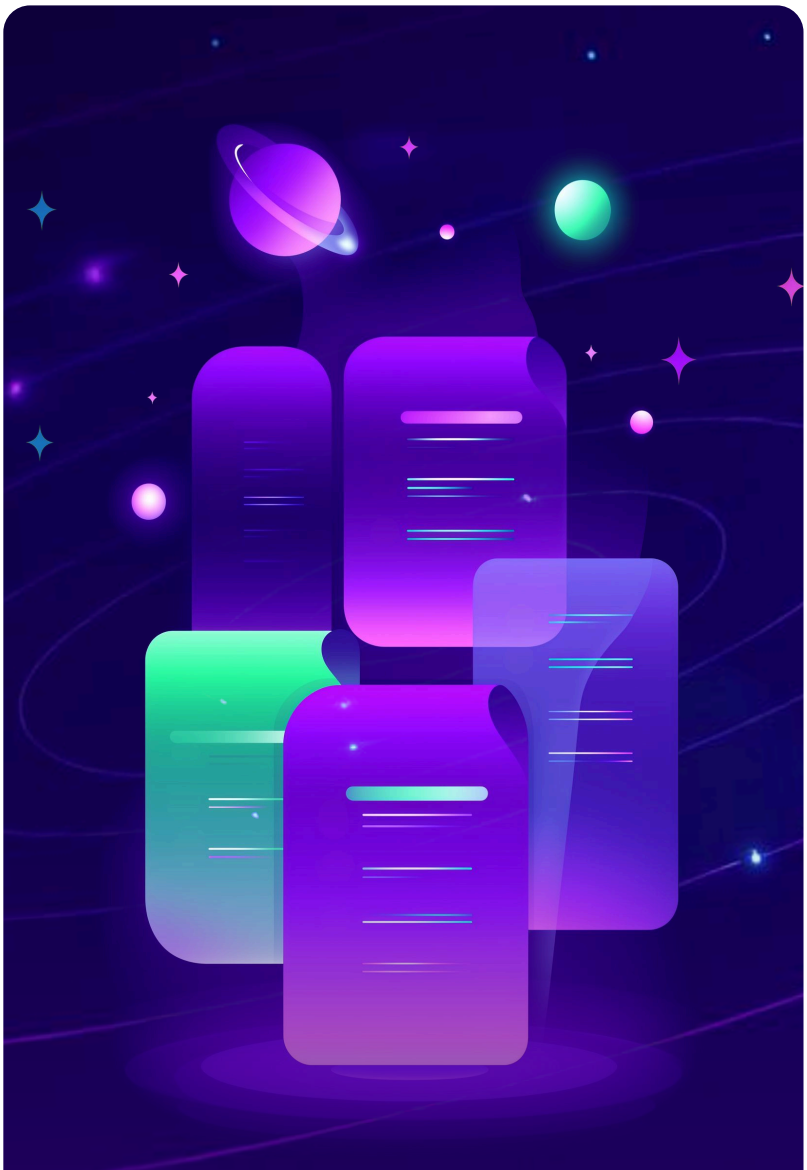

How Cyera Helps

Cyera AI Guardian inventories copilots and copilots-enabled users, tracks which data those tools can access, and classifies the sensitivity of that data. It inspects prompts and responses in real time, flags content that violates policy, and integrates with DLP tools to support enforcement.

Your internal copilots just got smarter - and hungrier for data.

Enterprise GenAI tools like Microsoft 365 Copilot, Gemini for Workspace, and ChatGPT Enterprise are embedded across everyday workflows - and they're drawing from unstructured, often poorly governed data. With default or overly broad access, these tools surface sensitive information fast, across teams and geographies, revealing long-standing blind spots in data governance.

What used to be slow or isolated oversharing now scales instantly. When copilots use retrieval-augmented generation (RAG) to comb through organizational content, the risk isn't just exposure - it's amplification.

Why it matters

Most teams can't yet answer three basics: What can each copilot see? Who and what is actually using them? And should they?

What good looks like

- Inventory of all enabled copilots, users, and service accounts
- Continuous discovery and classification of all data copilots might access
- Dynamic policies that auto-adjust based on data sensitivity and user role

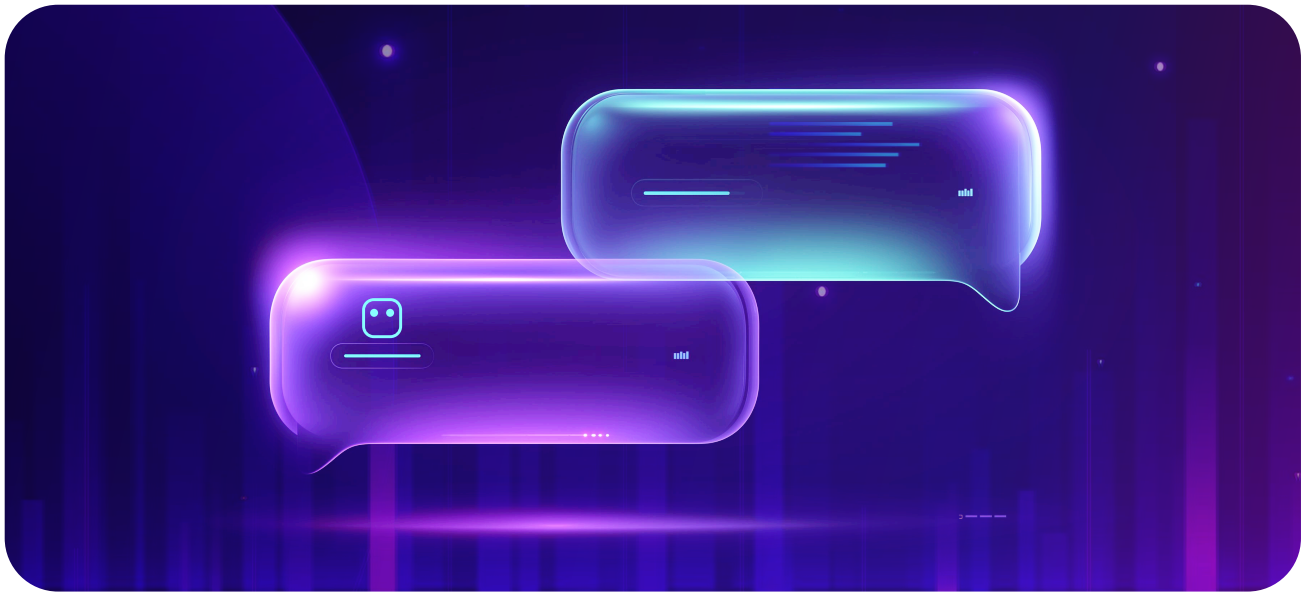

AI prompts are the new copy-paste.

Developers paste code for debugging. Product teams upload strategy decks. Customer support feeds in transcripts. Because ChatGPT is just that fast. But ease equals exposure. In a 2024 ESAF survey, 61 percent of Fortune 1000 CISOs cited GenAI-driven intellectual property leakage their top concern.

Why it matters

When data is sent to a public LLM, security teams lose visibility over how it's stored, retained, and reused. That undermines auditability, policy enforcement, and IP protections.

What good looks like

- Consistent policy enforcement across public AI tools
- Prompt-level logging from input to response
- Real-time detection and analysis that enables coaching

How Cyera Helps

AI Guardian logs interactions with public LLMs like ChatGPT and Claude, classifies the data involved, and monitors for sensitive information being shared. It tracks user and machine activity and supports blocking and coaching via DLP integrations when policy violations are detected.

Source: ESAF, How Top CISOs Are Navigating GenAI Risk & Opportunity, 2025

04 Homegrown AI

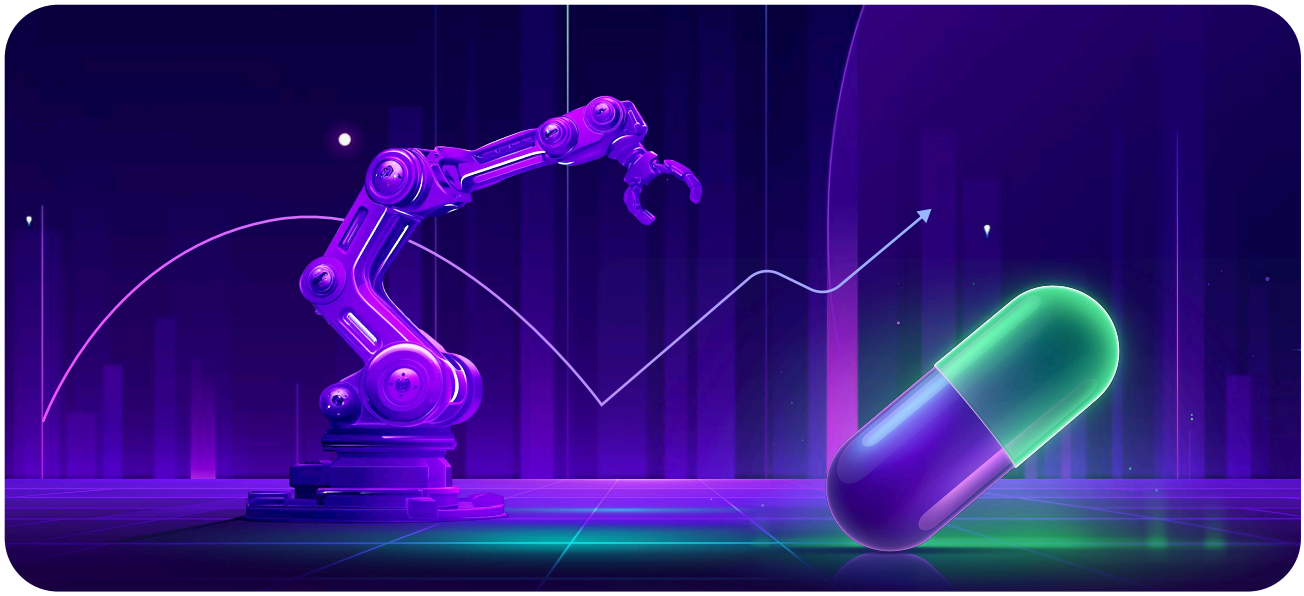

Build fast. Govern smarter.

Homegrown AI models fine-tuned on proprietary data are unlocking domain-specific insights and competitive advantage. But as these models grow more autonomous, governance becomes essential for continuous oversight and risk mitigation. According to Accenture, 3x more organizations plan to invest in agentic architectures in 2025 than in 2024, ushering in a new era of fast-moving, embedded AI.

Why it matters

Organizations that successfully scale domain-specific GenAI applications are more likely to outperform on ROI. But that value only materializes when data protection is built in from the very beginning.

What good looks like

- Automate the discovery and classification of sensitive data within training datasets.
- Apply runtime data protection to restrict content to only what each user is authorized.

How Cyera Helps

Cyera inspects prompts and responses in real time, flags sensitive data, and restricts the ingestion of sensitive or regulated content in homegrown apps. Your internal models stay fast, compliant, and secure - without slowing teams down.

05 AI Compliance

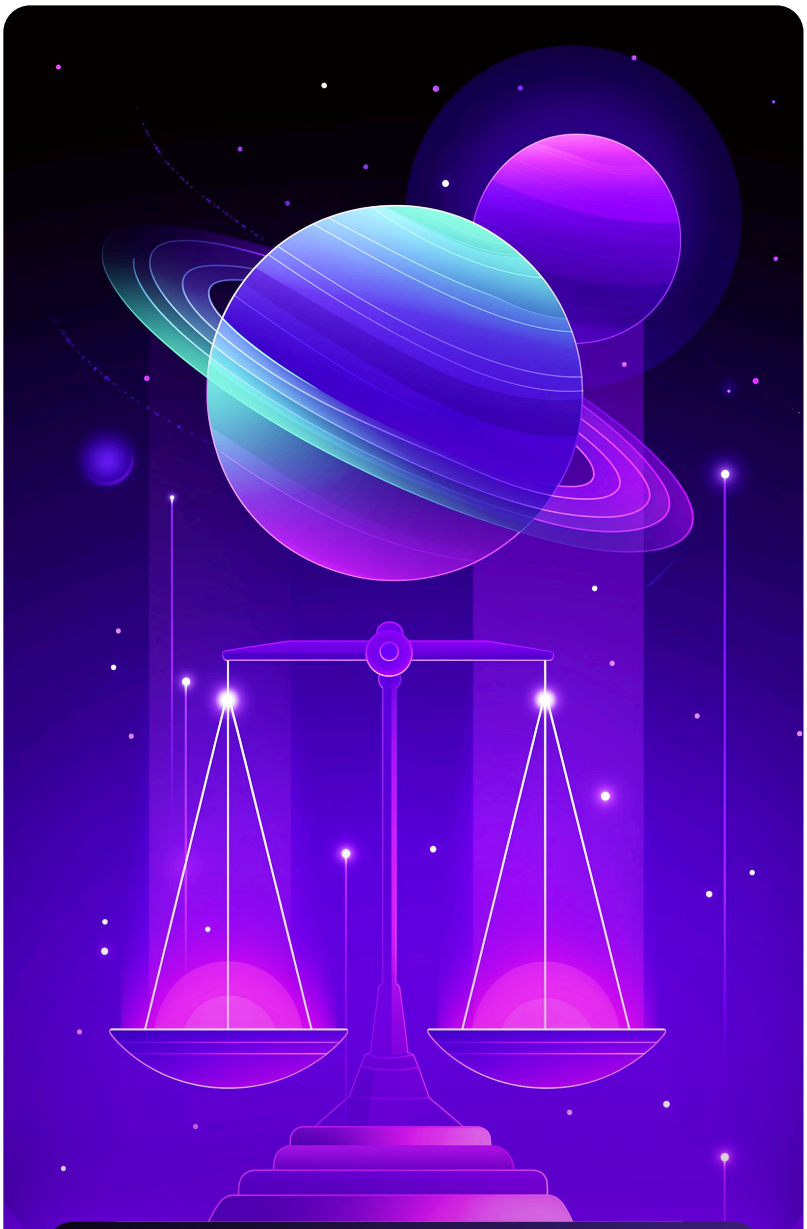

How Cyera Helps

AI Guardian inventories all AI systems, flags unauthorized access to regulated data, and classifies the information those models interact with. It supports enforcement of data minimization policies by helping teams identify and reduce access to stale, overexposed, or non-compliant data.

> Regulation isn't coming. It's already here.

Governments are moving fast to catch up with AI. In 2024 alone, U.S. federal agencies introduced 59 AI-related regulations - more than twice the year before. Globally, countries are shifting from pilot programs to enforcement, reshaping how enterprises approach governance.

Sixty-five percent of organizations say GDPR is influencing how they use AI, while 41% cite the EU AI Act.

And it's not just about compliance - it's about staying competitive. Nations are investing billions to scale AI safely, making secure, well-governed adoption a strategic advantage.

Why it matters

AI systems process massive volumes of sensitive data, making data minimization a regulatory and operational imperative. To comply, organizations must prove they're collecting only what's needed, retaining it only as long as necessary, and protecting it at every step.

What good looks like

- Define and enforce data minimization policies across AI workflows
- Track which models, agents, or apps are accessing sensitive data
- Conduct risk assessments tied to evolving compliance frameworks

Securing AI Starts with Your Data

AI systems are interacting with sensitive enterprise data every day - often without consistent oversight. That opens the door to data leakage, policy violations, and compliance risk. Security and data leaders need deep insight into how AI is using their data: what it can access, who's using it, and whether usage aligns with policy.

Cyera takes a data-first approach to AI security with **AI Guardian**, which provides **AI Security Posture Management (AI-SPM)** and **AI Runtime Protection**. AI-SPM inventories AI systems, users, and the data they access. Runtime Protection inspects prompts and responses in real time, flags risky activity, and enforces policies where needed. AI Guardian helps organizations discover where AI is running, understand what data is being accessed and by whom, detect risky behavior like prompt injection or unauthorized use, and support compliance with evolving regulations such as the EU AI Act.

AI won't slow down. But it can be governed. And it starts with knowing exactly what it's doing with your data.

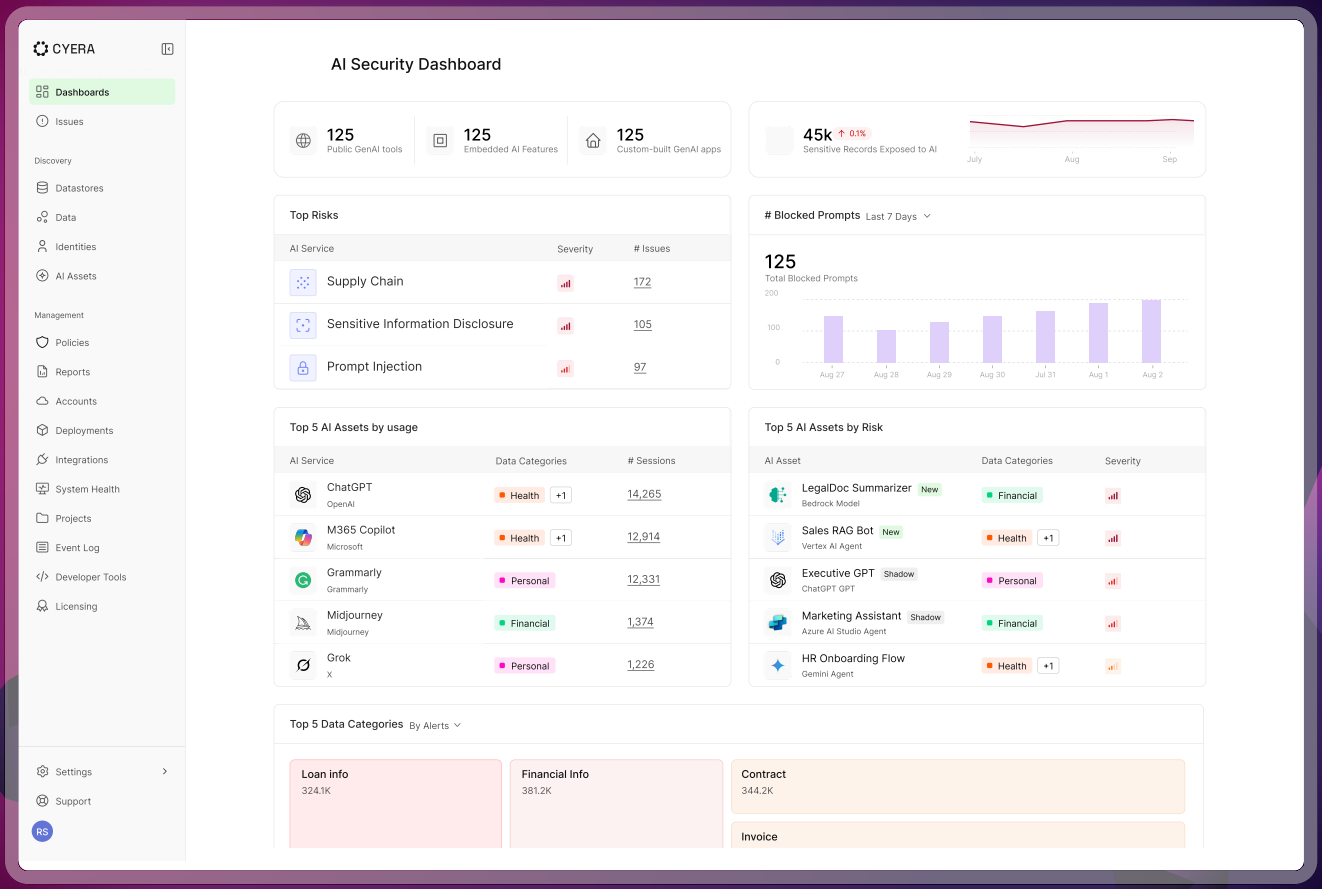

CYERA

A Unified Platform for Everything Data Security

From ground to cloud, Cyera powers your data security mission with a single platform designed to eliminate your most critical data risks - efficiently and effectively.

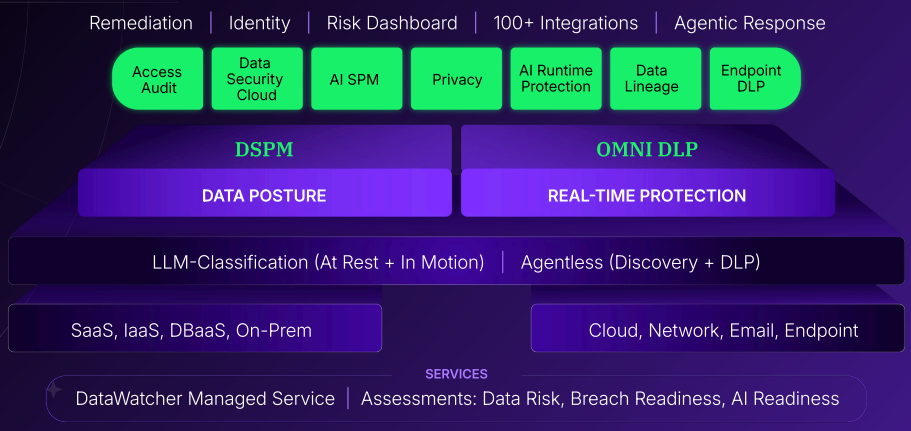

Data Security Posture Management (DSPM)

Discover and classify sensitive data across the cloud and on-prem with speed, scale, and precision. Remediate data vulnerabilities, reduce your attack surface, and realize cost savings from day one.

AI Security Posture Management (AI SPM)

Uncover every AI tool, surface data risks related to AI usage, and enforce least privilege policies to ensure AI only touches what it should.

Access Audit

Track who accessed what data and when to facilitate forensic investigations, identify insider threats, and discover unused sensitive data.

Privacy

Build an inventory of personal data including PII, PHI, and PCI. Identify privacy risks, streamline SAR responses, and demonstrate compliance with privacy regulations.

Omni Data Loss Prevention (Omni DLP)

Discover, classify, and protect sensitive data across the enterprise in real-time. Automate policy enforcement, minimize false positives and alert fatigue, and maximize visibility and control at scale.

AI Runtime Protection

Continuously inspect prompts and responses, detect threats like prompt injection or policy violations, and automatically block risky AI behavior before it causes harm.

Data Lineage

Trace your data throughout its lifecycle to help you build better DLP policies, comply with audit requirements, and remediate vulnerabilities at their source..

Data Security Cloud

Have a conversation with your data to prioritize risk remediation and glean insights to drive your business.

Why Cyera?

- Agentless Discovery, Faster Data Visibility
- AI-Native Classification with 95%+ Precision
- Automated Risk Identification and Prioritization
- AI-Native Brain to Bring Together Your DLP Stack
- Unmatched Scale to Match Your Data Growth

For more information visit cyera.com or scan→