Lightbeam

SANS | Research Program

Finding Sensitive Data Isn't the Hard Part: How Lightbeam Answers Whose Data AI Can Reach, and What to Do About It

Insights from the 2026 SANS AI in Cybersecurity Survey

PRODUCT BRIEF

The 2026 SANS AI in Cybersecurity Survey describes a governance function expanding faster than its own visibility. 76% of cybersecurity teams now hold a role in governing enterprise AI, but only 36% have a formal AI risk management program, and 63% cite lack of visibility into AI model use cases and risk exposure as a governance challenge. Underneath both numbers sits a more specific problem, 36% of practitioners name sensitive data or company IP leakage through employee use of AI tools as a top emerging threat. Lightbeam was built to answer the question underneath all three: not just what sensitive data exists, but whose data it is, who or what can reach it, and whether that access is appropriate.

Why Finding Sensitive Data Isn't the Same as Governing It

Each finding describes the same shortfall from a different angle. Teams are accountable for AI governance without the program to back it up, they lack visibility into what AI can reach, and the leakage they worry about most is already happening through ordinary employee use of AI tools. Knowing that sensitive data exists has never been the hard part.

Lightbeam's Data Identity Graph is built to close that gap. It connects data objects, data subject entities, accessor identities, permissions, AI interactions, and business context, so teams can move from "we found sensitive data" to "we know whose data is exposed, who can access it, whether AI is being used to expose or extract it, and what needs to change." That shift, from discovery to identity, is what separates a dashboard from a governance program that can produce audit-ready evidence.

Key Findings

76%

Governance roles are outpacing governance programs.

76% of teams now hold a governance role, but only 36% have a formal AI risk program to back it up. Lightbeam closes that gap by mapping sensitive data to the people and entities it represents, giving governance a concrete, auditable starting point.

53%

Lack of visibility starts at the data layer.

53% of practitioners call this their top governance challenge, and most of it traces back to not knowing what sensitive data an AI tool can reach. Lightbeam's Data Identity Graph connects data, identities, and AI interactions so teams know exactly what's exposed.

36%

Employee AI use is already leaking sensitive data.

36% of practitioners name sensitive data or IP leakage through employee AI use as a top emerging threat, already happening through sanctioned tools and shadow AI alike. Lightbeam monitors prompts, responses, and outputs, tracing any exposure back to the user, data subject, and source file.

Platform Capabilities

Lightbeam is built for security leaders at mid-to-large enterprises in regulated or data-intensive industries adopting Microsoft Copilot, Gemini, ChatGPT Enterprise, Claude, SaaS AI, or agentic AI workflows across complex, multi-platform data environments (see Figure 1).

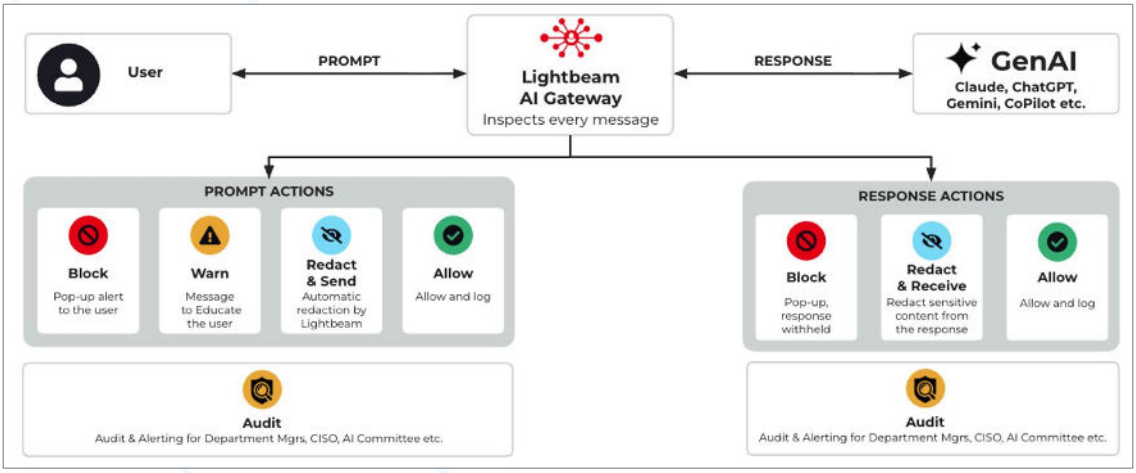

Figure 1. Lightbeam Data Identity Graph and AI Exposure Interface

Stop Shadow AI

Unauthorized AI use often puts sensitive data outside the organization's control before anyone notices it left. Lightbeam helps teams discover and control this activity by monitoring AI traffic across supported GenAI tools and inspecting prompts, responses, uploaded files, referenced content, and attachments. This visibility allows teams to identify when employees are using unsanctioned AI tools with regulated, confidential, or proprietary data. As policies mature, teams can block unauthorized AI use, warn users, or allow and audit activity.

AI Data Security

Lightbeam discovers and classifies sensitive data across SaaS, cloud, collaboration platforms, databases, and on-prem systems, then maps it through the Data Identity Graph to the people, customers, employees, vendors, business entities, and intellectual property it represents. Lightbeam's AI engine correlates AI interactions with data already scanned across the enterprise, distinguishing real production data from synthetic, masked, or test data so AI policies can decide what should be allowed, restricted, redacted, or blocked.

Least Privilege AI Access

Excessive access for users, applications, copilots, agents, and third-party SaaS AI creates risk that AI can quickly amplify. Lightbeam reduces this exposure by connecting sensitive data to data identity, access rights, business context, and authoritative sources of truth, showing which access paths create the most risk so teams can remediate them. Policies built on this foundation can revoke risky access, tighten sharing, enforce least privilege, and limit what AI systems can retrieve, use, or expose.

AI Governance

Continuous enforcement and audit-ready evidence replace one-time assessments and dashboard-only visibility when AI governance runs through Lightbeam. Teams can assess AI data risk, enforce AI data security policies, minimize unnecessary AI-accessible data, and document every policy decision and remediation action. The result is a clear record of what happened, whose data was involved, what action was taken, and how risk was reduced.

Request a Demo

SANS Research Program