CUSTOMER STORIES

Securing AI Usage at a Global Investment Bank



Industry	Employees	Use Case
Financial Services	1,500	AI and data security governance

The Challenge

A global investment bank was concerned about safely enabling AI adoption after discovering that sensitive content had been shared with ChatGPT. They also needed a tool that could help them understand and protect the large volumes of sensitive and highly regulated data they manage for their corporate clients.

- **Lacked visibility into their data.** Needed help understanding their Box, Snowflake, and Microsoft 365 (OneDrive, Exchange, SharePoint) data, including which users had access, what they were doing with it, and the associated risks; also wanted help understanding how much stale and duplicate data they had.
- **No AI guardrails for enterprise data.** Concerned about Microsoft Copilot, ChatGPT Enterprise, and Claude Enterprise accessing sensitive data in their Box, Snowflake, and Microsoft 365 environments; needed to set up entitlement guardrails to understand their data permissions and classifications and secure the data.
- **Needed to monitor enterprise-sanctioned use of AI.** Wanted a solution that could actively monitor what was being shared with ChatGPT Enterprise and Claude Enterprise, protect their data, and demonstrate compliance; specifically required the solution to integrate with ChatGPT Compliance API and Anthropic Compliance API for Claude.
- **Lacked visibility into shadow AI.** Did not know what unsanctioned public AI tools were being used by their employees and what data was being input.

The Solution

Together, Concentric AI's AI and data security governance solutions, Semantic Intelligence™ and Semantic DLP, gave the bank's security leadership the visibility and deep data context they needed to understand their data estate. Powered by Concentric AI's data categorization, they were able to ensure highly sensitive M&A documents weren't shared with public AI and ChatGPT Enterprise—securing AI usage across multiple browsers, including Island, remediating sensitive data at risk to ChatGPT Enterprise access, and meeting organizational privacy and compliance requirements with daily reporting on user activity.

Results

Secure deployment of enterprise-sanctioned AI.
The customer successfully deployed Microsoft Copilot, ChatGPT Enterprise, and Claude Enterprise with full visibility into all user interactions. When sensitive content was shared, Semantic DLP warned, blocked, or redacted in real time. Integration with the Compliance APIs for ChatGPT and Claude gave the customer greater visibility into and control over AI usage.

Complete visibility into shadow AI.
Semantic DLP provided the customer with real-time visibility into all unsanctioned applications—the security team could see when employees were using personal ChatGPT accounts instead of ChatGPT Enterprise. The categorization provided by Semantic Intelligence helped them understand the types of data flowing in and out of these tools.

Category-aware DSPM and AI DLP.
The customer's data was accurately mapped into categories and sub-categories to ensure that M&A records and other sensitive data was not shared with their AI applications or public AI.

Meet internal requirements for logging enterprise and public AI usage.
The customer was able to export conversation histories and each day downloaded reports that detailed the complete text of every user prompt and the corresponding model response.

AI guardrails established for enterprise data.
Semantic Intelligence helped the customer understand what data was in Box, Microsoft, and Snowflake as well as the permissions that Microsoft Copilot and ChatGPT Enterprise were inheriting and using to access the data in these repositories.

Audit trail of data access and usage.
The data lineage feature in Semantic Intelligence gave the customer comprehensive insights into how data moved across the organization, who was accessing it, and what they were doing with it—every view, edit, download, and share.

Improved data hygiene.
The customer was able to find duplicate, near-duplicate, and stale data records and then easily remediate them in the platform.

Flexible path to operationalization.
The deployment was tailored to meet the customer's immediate concerns—addressing data risk in AI usage. Next, they worked with their partner and the Concentric AI customer success team to implement data classification and policy-defined metadata labeling.

Local deployment to meet privacy and compliance requirements.
The Semantic Intelligence platform was deployed within the customer's private Microsoft Azure cloud to meet regulatory requirements for processing highly sensitive data on premises.

Discover, monitor, and protect sensitive data everywhere AI is used.

Let's Talk ↘