SANS | Research Program

SPONSORED BY

bonfy.ai

Adapting Your Data Security Program in the Age of AI

A SANS First Look

Written by **Kevin Garvey** | January 2026

Introduction

Artificial intelligence is changing the way companies make major productivity and efficiency gains that support the business missions of their organizations. By harnessing internal data, organizations are unlocking tangible results as AI turns information into actionable insights. As with most technology revolutions, AI brings with it new cybersecurity concerns and challenges. Examples include:

- Expanded attack surfaces, created by additional integration points and data flows
- Rapid growth in AI-generated content that increases the risk of exposing sensitive information
- Immature or uneven data security programs struggling to keep pace with AI adoption
- Disparate tool sets unable to enforce consistent security policies across environments
- Shadow AI usage operating outside the visibility of traditional data loss prevention (DLP) solutions
- Existing data-labeling practices that are not mature enough to scale to AI inputs and outputs
- Inadvertent exposure of intellectual property, merger and acquisition data, or protected information such as personally identifiable information (PII) and protected health information (PHI)

Historically, traditional DLP and data security programs were created to manage legacy data at rest and in transit. Today, the internal data used as AI inputs and outputs is being bolted onto most organizations' existing DLP programs. Unfortunately, this patchwork integration into legacy DLP programs does not always protect data in the age of AI. Some areas where legacy DLP programs fall short in this area include:

- Legacy data security solutions are unable to understand the business context of AI-generated content.
- Programs are unable to identify all data in an organization, creating gaps in data governance.
- New datapoints are not governed, leading to unintended data security consequences.
- Data input into large language models (LLMs) could contain sensitive information, such as intellectual property or PII.
- DLP operational alerts do not have updated playbooks, leading to inadequate responses from a security operations team.
- DLP content tuning activities may not account for new data points or content generated by AI.
- An overreliance on pattern matching and keyword scanning overlooks "eyes on glass governance." This process requires constant and labor-intensive upkeep.
- User data governance may not be treated uniformly, leading to inconsistent application of data security governance.
- With a lack of visibility comes an ever-increasing threat of insider risk.
- The relationship among users, devices, AI agents, applications, data flows, and information may not be seen from a traditional program, leading to large governance inconsistencies.
- Data ingested into internal LLMs may contain information not meant to be shared with others inside the company, or—even more concerning—with others outside the company.
- Existing key risk indicators (KRIs) and key performance indicators (KPIs) may not have the appropriate scope and coverage to report and govern against AI data points. Additionally, the expanded scope of DLP metrics and dashboards may not consider new audiences using data from AI.

About Bonfy.AI

AI can bring in efficiency and productivity gains while maintaining a “security first” mentality. Bonfy.AI can be the foundation for a model data security program or augment an existing data security program (such as DLP) to comprehensively address risks of AI inputs and outputs. AI mandates shifting the focus of data security from being a reactive, contextless task into being a proactive, contextual enabler, allowing the enterprise to securely use AI without the cybersecurity concerns it introduces. The Bonfy Home screen is designed to give you a quick overview of your organization’s content-security posture and to let you jump into key workflows (see Figure 1).

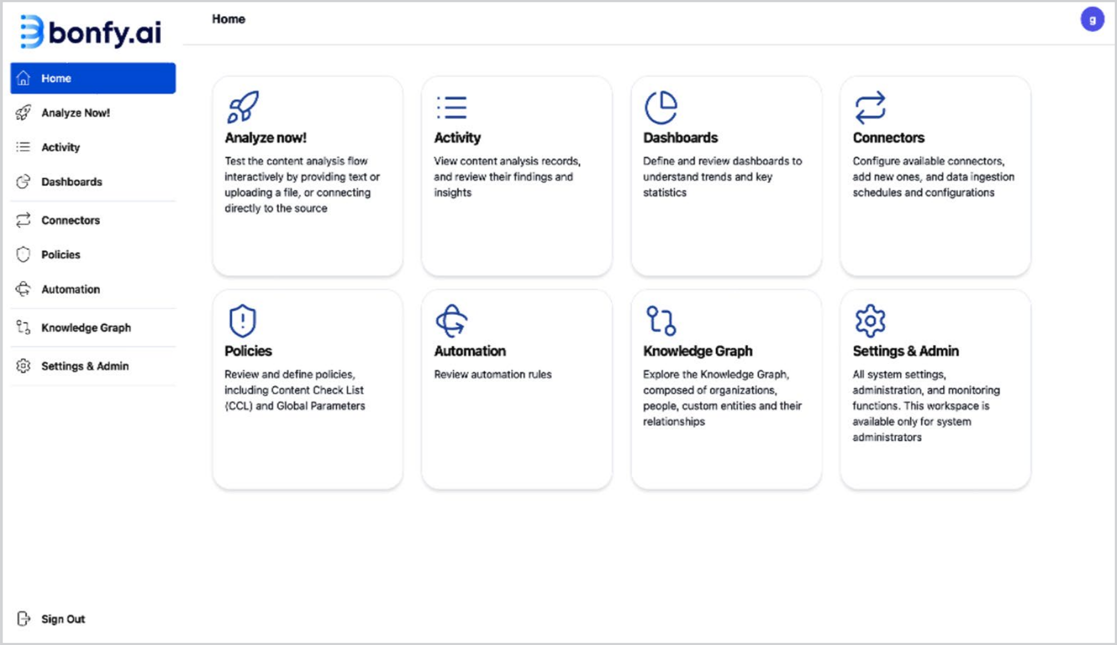

Analyze now!
Test the content analysis flow interactively by providing text or uploading a file, or connecting directly to the source

Activity
View content analysis records, and review their findings and insights

Dashboards
Define and review dashboards to understand trends and key statistics

Connectors
Configure available connectors, add new ones, and data ingestion schedules and configurations

Policies
Review and define policies, including Content Check List (CCL) and Global Parameters

Automation
Review automation rules

Knowledge Graph
Explore the Knowledge Graph, composed of organizations, people, custom entities and their relationships

Settings & Admin
All system settings, administration, and monitoring functions. This workspace is available only for system administrators

Figure 1. Bonfy Home Screen

Bonfy understands what data may be risky to an organization. Importantly, it continually learns about an organization’s data, so it can provide strong protection to the ever-changing data landscape in the context of the business and relevant entities. This is done by gathering and correlating data from different enterprise datastores, such as customer relationship management (CRM), identity and access management (IAM), and human resource management (HRM) tools, allowing Bonfy to create custom knowledge graphs that show what data may be at risk and from what entity (see Figure 2). Continuous harmonization allows Bonfy to quickly learn and adapt to changes in data and data strategy, resolving AI-based data security concerns quickly.

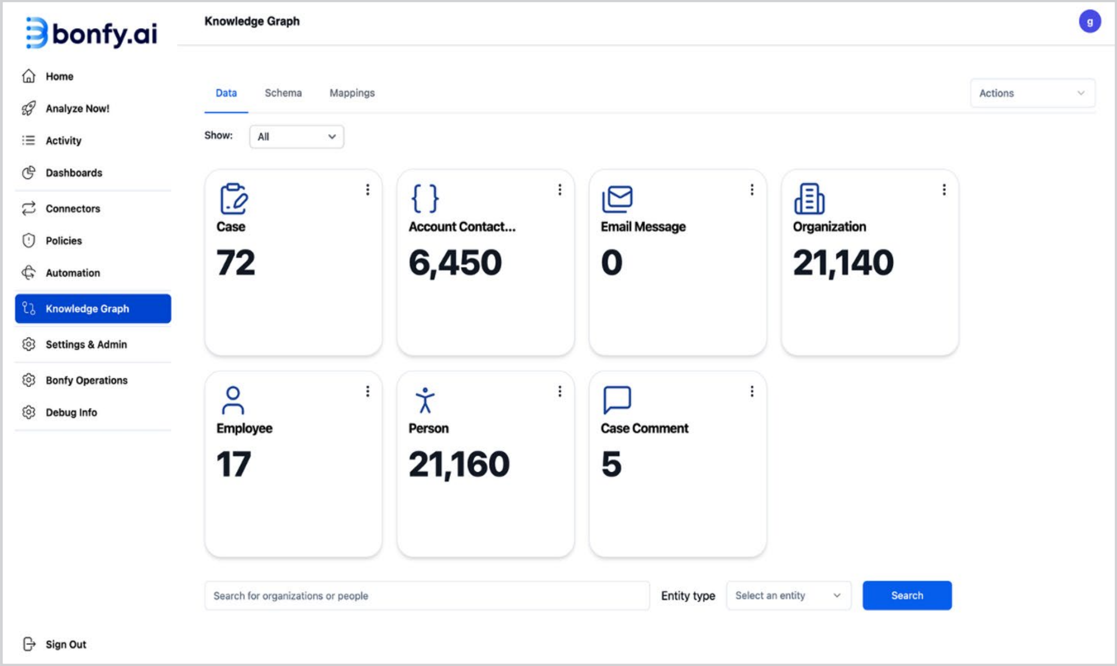

Figure 2. Bonfy Knowledge Graph

Bonfy can identify both upstream and downstream risks of any information flow, enabling proper protection for sanctioned AI-based applications (consumed or developed) and unsanctioned applications (Shadow AI).

Importantly, traditional data security programs are labor-intensive, often causing delayed responses when new measures are needed. Bonfy allows enterprises to move beyond inefficient manual data security programs of the past to a dynamic way of protecting changing data strategies in the new AI world. Leveraging its multichannel architecture, Bonfy allows for the consolidation and sunsetting of old tools and outdated processes in your organization.

The manual days of needing to label data are over with Bonfy. Its entity-aware engine automates data labeling and classification with a streamlined approach that is adapted to the specific organization to manage data at rest and in transit. Labeling is riddled with tool set and user errors in most organizations, so automated data labeling is a huge win for any data security program.

Lastly, legacy data security program management using non-AI-based security tools is difficult to measure and report. Bonfy solves the management of data security governance and program management by providing a set of easy-to-use out-of-the-box dashboards, policies, and metrics. Near-real-time data security metrics and reporting can be customized to fit the needs of all stakeholders in the organization, such as risk and compliance teams.

Conclusion

Bonfy enables enterprises to pivot from a legacy data security strategy to AI-enabled data security by allowing them to:

- Dynamically scale data security programs as the enterprise's data needs grow
- Automate data security and governance management by enabling executives to showcase the detection and prevention of data risk
- Ensure that data to and from AI-based deployments is being secured quickly
- Offer data security leaders an opportunity to rethink and redeploy nimbler, more efficient, and more resilient data security programs

Bonfy ensures a comprehensive and adaptable view of enterprise data, secures the data to and from AI applications, and dynamically scales to the growing data needs of the enterprise. To remove the stress of constant AI data security headaches in the enterprise, consider leveraging Bonfy.