BONFY SOLUTION BRIEF:

Data Security for the AI Era

bonfy.ai

Organizations are trying to accelerate AI adoption across Copilot, SaaS applications, collaboration platforms, and emerging AI agents, but most data security programs were built for static channels and human-paced workflows rather than always-on, multi-step AI interactions. Bonfy addresses this gap with a unified, entity-aware data security platform that protects unstructured data across email, files, SaaS apps, collaboration tools, AI systems, and agentic workflows.

When AI Acceleration Outruns Data Security

As AI becomes embedded in business workflows, sensitive content moves through more channels, more tools, and more autonomous processes, increasing the risk of oversharing, privacy violations, integrity issues, and compliance failures. Traditional DLP and DSPM tools often lack the contextual understanding, cross-channel visibility, and accuracy required to enforce policy safely in these environments, which leads many organizations to narrow controls, rely on monitoring, or avoid enforcement altogether.

Why legacy controls fall short

- Analyze content in isolation
- Depend on static rules or shallow classifiers
- Struggle to distinguish generic content from customer-, employee-, or consumer-specific information.

The result:

These limitations matter more in AI workflows because enforcement must increasingly happen in line, not after the fact, and in-line enforcement only works when systems understand the business context well enough to make accurate decisions.

Bonfy ACS: For Data Security

Bonfy provides adaptive content security built around contextual, entity-aware analysis so organizations can detect, classify, label, investigate, and prevent data risk with greater precision across both human and AI-driven workflows. Its platform applies one intelligence layer and one policy framework across data in motion, at rest, and in use, reducing blind spots and tool sprawl while helping teams move from visibility to automation and prevention.

How Bonfy works

Bonfy's analysis engine combines insight extraction, entity enrichment, policy evaluation, risk scoring, tagging, and explainability to understand not just what content says, but who it relates to and why it matters in context. That intelligence supports real-time and near-real-time controls across outbound communication, collaboration systems, file repositories, SaaS applications, and AI-enabled workflows, including environments where latency matters for browser, API, or agent interactions.

Key capabilities

- Multi-channel protection across email, files, SaaS apps, collaboration tools, AI systems, and AI agents.
- Entity-aware classification and detection that reduces false positives and enables safer enforcement.
- Unified policy and automation engine for discovery, labeling, remediation, and enforcement.
- Automated contextual labeling, including support for Microsoft Purview sensitivity labels.
- Response orchestration and channel-specific prevention actions such as block, quarantine, redirect, modify, and notify.
- Entity risk management for humans, systems, and AI agents, with risk attribution and historical insight.
- Cloud-native deployment with SaaS and bring-your-own-cloud options.

AI and agent security

Bonfy's data security approach for AI extends across three control layers: controlling what data AI systems and agents can access, inspecting what they produce before it leaves through downstream channels, and enabling in-process inspection during agent reasoning through Bonfy's MCP server capability. This positions Bonfy around the data layer of AI security rather than only configuration management, helping organizations secure prompts, retrieved content, tool calls, external MCP interactions, embeddings, and outputs with the same platform intelligence.

Where Bonfy delivers value

USE CASE	CUSTOMER PROBLEM	BONFY VALUE
Outbound email and collaboration	Sensitive information leaks through email, file sharing, and collaboration tools, while legacy DLP creates noise and misses business context.	Bonfy analyzes content with context and entity awareness to improve detection accuracy and enable policy-aligned prevention.
Copilot and AI app security	Organizations lack visibility into what AI tools can access, index, or generate.	Bonfy provides upstream and downstream visibility and enforcement so sensitive content does not flow into the wrong prompts, indexes, or outputs.
CRM, ITSM, and SaaS sprawl	PII, secrets, and customer data accumulate across systems like Salesforce, HubSpot, Slack, Jira, and ServiceNow.	Bonfy discovers and governs sensitive content in operational systems with high accuracy and lower operational friction.
AI agent workflows	Agents can access data, call tools, and take actions without human review, creating new leakage paths.	Bonfy secures the data lifecycle around agents by controlling grounding, inspecting data in use, and checking downstream outputs.

Outcomes

With Bonfy ACS, organizations can:

- Reduce blind spots across the modern data surface
- Improve trust in enforcement
- Support AI adoption without depending on multiple disconnected tools
- Gives teams a phased path from visibility to prevention
- Fast time to value and lower operational burden for already-stretched security programs

Bonfy helps organizations protect sensitive unstructured data everywhere it moves, lives, and is used, across human workflows and AI-driven workflows alike.

By combining multi-channel coverage, contextual accuracy, and unified enforcement, Bonfy turns data security from fragmented detection into practical control for the AI era.

Contact Us

For inquiries or additional information, please contact us at info@bonfy.ai

bonfy.ai

/bonfy-ai

bonfy.bsky.social