SOLUTION BRIEF

Bonfy Contextual Data Enforcement

Bridging Enterprise AI Adoption and Data Governance

bonfy.ai

bonfy.ai

The Challenge:

AI Adoption Without Data Control

Enterprises are racing to connect AI systems (Claude, Microsoft Copilot, ChatGPT Enterprise, Perplexity, and others) directly to their data stores, such as SharePoint, Microsoft 365, Google Drive, and beyond. The productivity benefits are undeniable. The security risks are equally real.

Today's native AI connectors authenticate on behalf of the user and expose everything that the user can access, with no content-level filtering, no policy enforcement, and no visibility. Organizations face an impossible choice: unlock the full productivity of AI or protect their sensitive data.

This gap is uniquely dangerous. Organizations cannot use Microsoft Purview sensitivity labels to control what a third-party AI like Claude can access. They cannot prevent PII, financial records, legal documents, or client-confidential data from flowing freely into AI context windows. Security and compliance teams are flying blind.

The Solution:

Contextual Enforcement Layer

Bonfy's Contextual Data Enforcement is a transparent enforcement proxy that sits between any AI client (Claude, Copilot Studio, and others) and enterprise data repositories (Microsoft 365, SharePoint, Google Drive). It leverages Bonfy's existing content inspection engine to review data in real time, before it reaches the AI, and enforces contextual data access policies at the content level, not just the user level.

How It Works

- Users or administrators replace the native AI connector (e.g., Claude's built-in Microsoft 365 connector) with Bonfy's secure connector, a simple configuration change.
- Traffic flows through Bonfy's enforcement layer, which calls the same underlying Microsoft Graph API or Google Drive API on behalf of the user.

- Bonfy inspects content in transit (detecting PII, sensitive classifications, confidential labels, and policy-defined categories) using the same analytics engine already deployed in the platform.
- Based on policy, Bonfy allows, blocks, or redacts data before returning it to the AI client. The AI receives a clean, policy-compliant response.
- User authentication and access permissions remain fully intact, Bonfy adds a layer on top, never replaces existing controls.

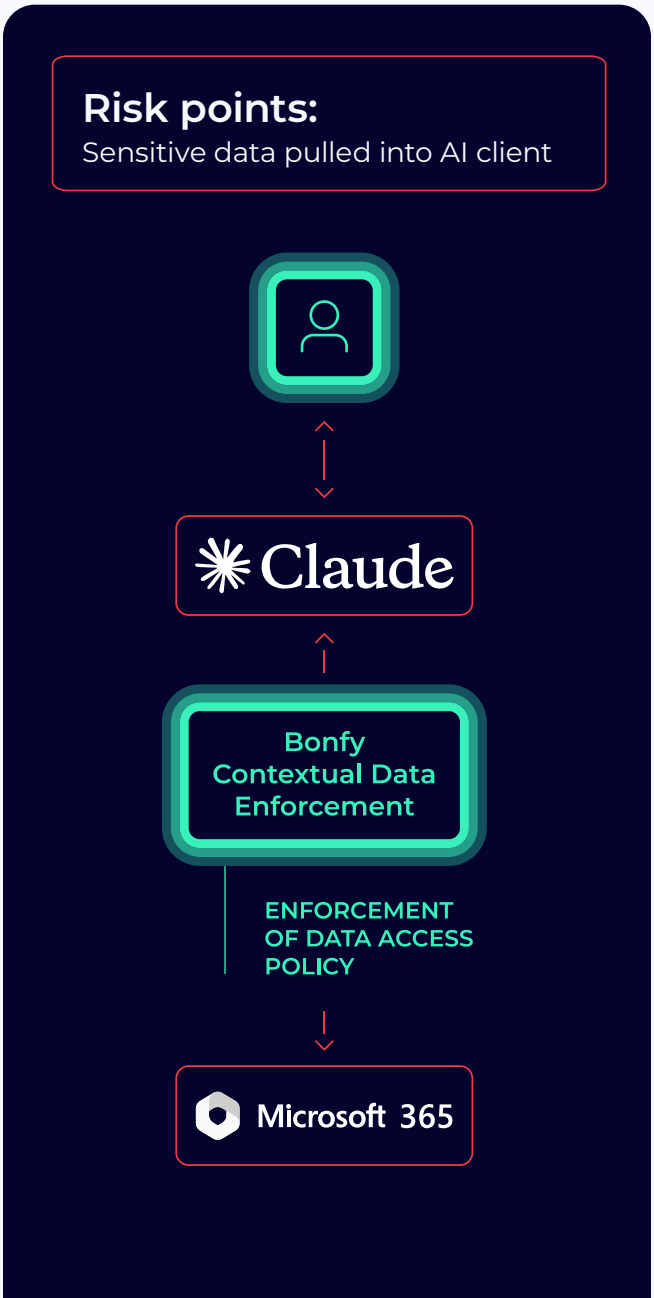

Key Capabilities

- **Content-Level Policy Enforcement:** Block or allow data based on content type, sensitivity classification, PII detection, and custom policy rules, not just file-level permissions.
- **Transparent Proxy Architecture:** No changes to the end user's workflow. The Bonfy connector looks and behaves identically to the native connector from the AI client's perspective.
- **Multi-Platform Support:** Works with Claude (Claude.ai, Claude Code, Claude Enterprise), Microsoft Copilot Studio, and any AI platform communicating via MCP or REST API.
- **Multi-Repository Coverage:** Microsoft 365 / SharePoint available at launch; Google Drive and additional repositories on roadmap.
- **Full Coexistence:** Works alongside existing CASB deployments, MCP gateways, and Purview policies, complementing, not replacing, existing investments.
- **Capability of Bonfy Engine:** Built on Bonfy's proven data classification, scanning, and inspection infrastructure — no new algorithms, no new risk surface.

Why Bonfy — Why Now

The window for securing enterprise AI adoption is open today. Organizations are deploying AI connectors to production environments without adequate controls. The options currently on the market are either overbuilt, dependent on complex infrastructure, or not purpose-built for AI data access. Bonfy's solution is purpose-built, lightweight, and deploys in the time it takes to switch a connector. It requires no new infrastructure, no gateway procurement, and no dependency on third-party MCP services. For security teams being pressured to enable AI adoption while maintaining compliance, it provides a simple, defensible answer.

Business Value

- Accelerate AI adoption safely — remove the primary security inhibitor without blocking productivity.
- Demonstrate compliance to clients and auditors — provide evidence of technical data controls, not just policy documents.
- Reduce incident exposure — prevent sensitive data from silently entering AI context windows or training feedback loops.
- Simple procurement and deployment — land with a focused use case, expand to full Bonfy platform capabilities over time.

Contact Us

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

bonfy.ai

/bonfy-ai

bonfy.bsky.social