SOLUTION BRIEF

Shadow AI Data Protection

Bring shadow AI into the light — see every unsanctioned copy/paste and upload while stopping the ones that put your data at risk.

The Problem

Employees have already adopted AI with or without security's approval. They paste source code into ChatGPT to debug, upload customer lists into Claude to summarize, and prompt Copilot with contract language they shouldn't share. Each interaction moves sensitive data into models your organization doesn't own, control, or audit.

Traditional DLP was built for email attachments and file shares, not prompts and pastes into AI tools. It can't see the activity, let alone stop it. This leaves security teams blind to the fastest-growing exfiltration channel in the enterprise.

The Solution

Mimecast Incydr gives security teams complete visibility into shadow AI activity for every paste, upload, and file transfer across unsanctioned AI tools. It includes adaptive controls to educate, block, or contain risky behavior without slowing the business down.

Incydr deploys in days, not months. No regex policies. No content classification project. On day one, you see which employees are using which AI tools, what data they're sending, and where the real risk lives.

Why Mimecast Incydr for Shadow AI

- **Day-one visibility** into every paste and upload to unsanctioned AI tools.
- **No policy lift required:** Incydr automatically tracks all unsanctioned data movement and prioritizes risk from day one.
- **File-level detail** on what's being exfiltrated: Source code, customer data, and IP.
- **Adaptive controls** from in-the-moment coaching to hard blocks that are right-sized to the user and the data.

10,000+ data exfiltration attempts to unsanctioned AI tools every hour across Incydr customers	**43%** of data exfiltrated to shadow AI is source code — among the highest-value assets in any organization	**80%** of cybersecurity leaders worry employees are concerned about sensitive data leaks through GenAI tools
Anonymized Mimecast Incydr Data, 2026	*Anonymized Mimecast Incydr Data, 2026*	*Anonymized Mimecast Incydr Data, 2026*

Why shadow AI breaks traditional data protection

GenAI adoption has outpaced every security category built to contain it. Three structural gaps make legacy DLP, CASB, and UEBA inadequate for shadow AI:

- **Prompts aren't files.** Employees exfiltrate data by pasting it into a browser, not by emailing it or uploading an attachment. Content-scanning DLP never sees the text.
- **The tool list changes weekly.** Employees try a new AI app every Monday. Solutions that rely on allowlists or pre-built connectors fall behind the moment a new tool gets traction.
- **Intent matters more than content.** A marketer pasting a press release into an AI summarizer is fine. An engineer pasting unreleased code into the same tool is a crisis. Content alone can't tell the difference — user, file source, and destination do.

The result: security teams know shadow AI is happening, but can't see who, what, or how much. Policy stays theoretical. Acceptable Use is unenforced. And the clock is ticking on the first breach that traces back to an AI prompt.

How Incydr protects data in the age of shadow AI

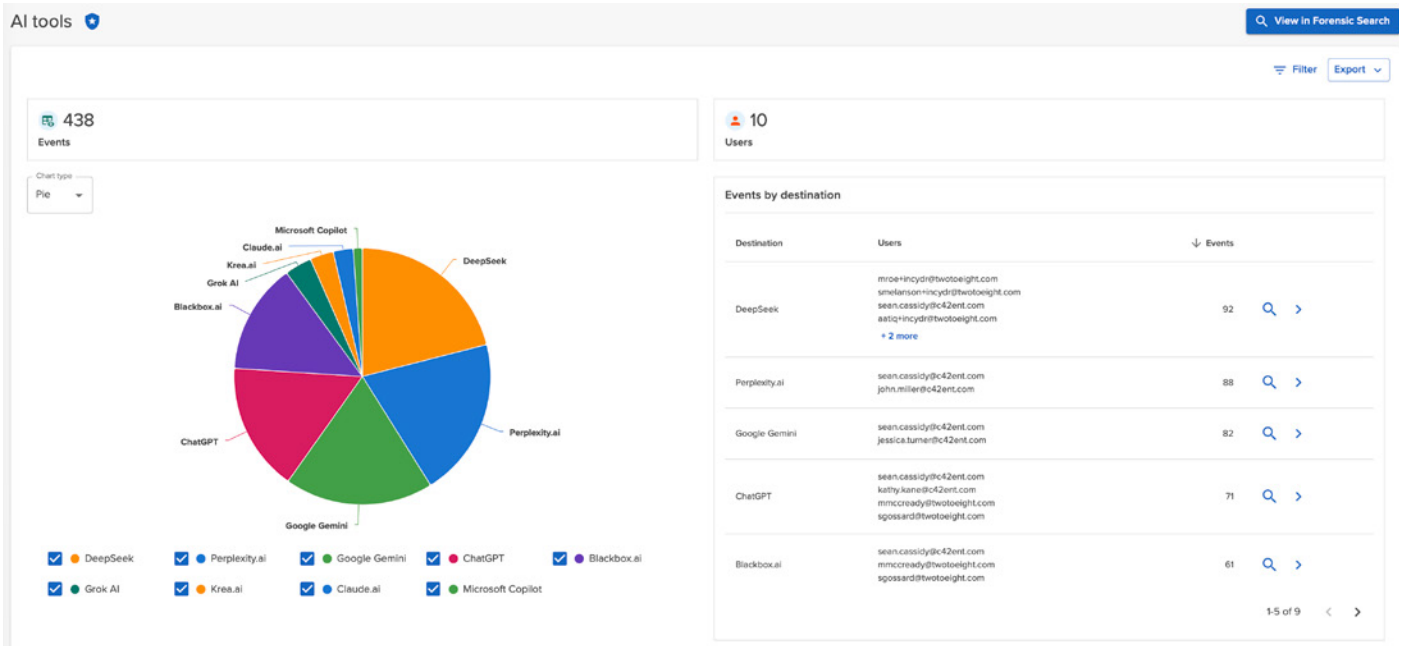

Incydr was built for the way employees actually work today across browsers, endpoints, clouds, and AI tools that didn't exist a week ago. It pairs day-one visibility with adaptive controls that scale to the risk, not the user's job title or the compliance checklist.

See every shadow AI interaction

- **Browser-based monitoring** captures unsanctioned copy/paste and file uploads across every AI destination without requiring pre-built connectors or site allowlists.
- **File-level visibility** shows exactly which document, snippet of source code, or record left the environment and who sent it.
- **Unsanctioned tool discovery** surfaces new and emerging AI applications employees adopt, so you can make informed policy decisions instead of reactive bans.
- **Customizable risk indicators** identify anomalies against normal activity, flagging the high-risk outliers against the noise of daily AI use.
- **Dedicated dashboards** for CISOs and analysts show shadow AI trends, top users, top destinations, and the sensitivity of data moving at a glance.

Control without stifling productivity

- **In-the-moment coaching** delivers short Instructor video lessons when a user pastes into an unsanctioned AI tool, teaching them which sanctioned alternative to use instead.
- **Block by source** prevents high-value data like source code, HR system exports, or CRM contents from reaching any untrusted AI or other destination.
- **Block by destination** stops uploads and pastes to specific unsanctioned AI domains for watchlisted users or tenant-wide.
- **Temporary allow** lets administrators allow employees to self-report a one-off need in the moment, while security retains a full audit trail.
- **Watchlist automation** applies stricter controls to departing employees, contractors, and other high-risk users while keeping the rest of the workforce unencumbered.

The shadow AI use case, start to finish

Day 1: Deploy the Incydr endpoint agent and browser extension. Within hours, dashboards populate with every unsanctioned paste and upload employees are sending to AI tools including the tools your security team didn't know were in use.

Week 1: Identify the top sanctioned and unsanctioned AI tools in your environment, the users driving the most risk, and the data types most frequently exfiltrated. Use built-in risk indicators to separate routine activity from genuine exposure.

Week 2: Turn on Instructor micro-trainings for low-risk events to coach behavior change in the moment. Event volume drops without a single hard block.

Month 1: Apply blocking for the scenarios that matter most: Source code leaving engineering, customer data going to any AI tool, or high-risk unsanctioned tools. Temporary Allow provides legitimate edge cases a safety valve with every justification logged.

Ongoing: Incydr flags and automatically adds new AI tools as employees adopt them. Your AI governance scales without scaling headcount.

Incydr Shadow AI capabilities at a glance

Capability	What Incydr delivers for shadow AI
Shadow AI discovery	• Browser-level capture of pastes and uploads to any web destination with no pre-built connector required, so new AI tools are visible the moment employees adopt them. • Automatic categorization of known GenAI destinations (ChatGPT, Claude, Gemini, Copilot, Perplexity, GenAI desktop apps, and hundreds more) with ongoing updates so you never miss a new tool. • Event volume, user count, and data type breakdowns by destination, surfaced in purpose-built dashboards.
Exfiltration detection	• GenAI, endpoint, cloud, browser, and email monitoring from day one with no policy setup and no laborious content tagging project. • File-level detail on what was sent: filename, source application, size, and sensitivity indicators. • Source code protection including comprehensive Git push/pull monitoring for engineering teams. • AI-based Content Inspection available for PII, PCI, and custom data entities.
Adaptive controls	• Paste and upload blocking to untrusted or specific AI destinations. • Blocking by source (e.g., CRM downloads, code repos, etc) • Block unsanctioned GenAI desktop apps • Available integrated Instructor micro-trainings: Scenario-based videos triggered in real time and sent via email or collaboration tools. • Tenant-wide or watchlist-based controls for departing employees, contractors, or high-risk users
Risk prioritization	• AI-powered risk prioritization out-of-the-box across three dimensions: File, user, and destination • Reduces alert noise so analysts focus on the small fraction of shadow AI activity that represents real risk. • CISO dashboards show shadow AI trends and adaptive control impact over time.

Stop shadow AI from becoming your next breach

Get a 30-day Proof of Value (POV) to see your organization's shadow AI activity within hours and how to take action to protect sensitive IP

mimecast.com/use-cases/shadow-ai/

About Mimecast

Secure human risk with a unified platform.

Mimecast's connected human risk management platform prevents sophisticated threats that target human error. By gaining visibility into human risk across your collaboration landscape, you can protect your organization, safeguard critical data, and actively engage employees to reduce risk and enhance productivity.