

# The New Agentic Intelligence Layer for Data Loss Prevention

Automate the tedious work. Multiply your security team's impact without multiplying headcount.

## DLP Security Dashboard

### Alerts by Severity




■ Critical  
■ High  
■ Medium  
■ Low

### Alerts by Channel



■ Email  
■ Cloud Drive  
■ Generative AI

### Alert Severity Over Time



■ Critical ■ High ■ Medium ■ Low ■ None

### Top 5 Destinations

| Application                        | Alerts |
|------------------------------------|--------|
| ChatGPT AI                         | 59     |
| Microsoft Sharepoint Cloud Storage | 53     |
| Yahoo Email Service                | 81     |
| AOL Email Service                  | 67     |
| Slack Software                     | 72     |

## It's no longer humanly possible to manage DLP.

Your team starts every morning the same way. Thirty thousand alerts a week. Three analysts. 99% never get looked at. An analyst spends an hour investigating an incident that turns out to be a false alarm. Meanwhile, a real insider threat sits buried as a low severity alert. Now employees are pasting sensitive data into AI tools your stack was never built to see. No amount of hiring can solve these problems.

Now imagine something different. One intelligence layer across the DLP stack you already run. Alerts arrive pre-analyzed. Policies write and improve themselves. Every analyst action trains the system to get smarter. One customer saw 16,000 events reduced to a handful.

That's Cyera Omni. The new agentic intelligence layer for DLP.

## Challenges for Security Teams



### Policy Headaches

DLP can feel like an endless loop of tuning, testing, and breaking things. One bad rule, and you're blocking legitimate work or worse, letting true positives slip through the cracks.



### False Positive Fatigue

If 95% of your alerts are low-level noise, your team will start ignoring all of them. That leaves real risks buried under the chaos.



### Siloed Tools, Siloed Teams

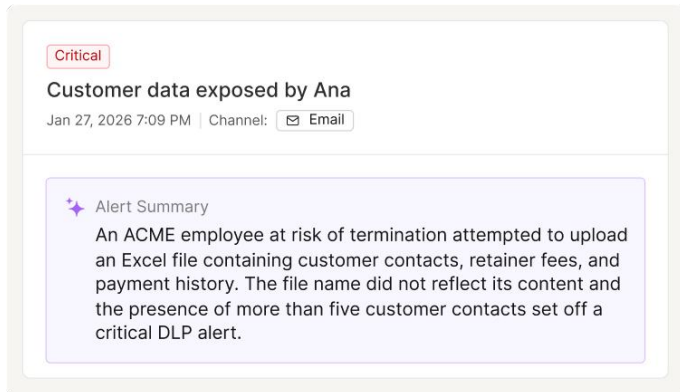
Cloud, network, and endpoint DLP each operate in isolation. Policy coordination becomes a time-sink. Visibility? Fragmented at best.



### Missing Context

Is that credit card number just Dave buying sneakers again, or an actual PCI violation? Without behavioral and data context, every alert looks equally urgent or risky.

# The Difference is Significant



## See the full story of insider risk.

Every alert arrives pre-analyzed and ready to act on. Go from investigation to triage in minutes, not hours.


- AI summaries explain what happened, what data was involved, where it was going, and why it matters
- Related alerts are grouped for deeper investigation
- Assign, adjust severity, update status, or inspect the raw data that triggered the alert.

## Skip the manual policy work.

Policies that write and improve themselves, closing coverage gaps as new behavior patterns and data flows emerge. Protection improves over time without heavy manual lifting by analysts and engineers.

- Your unique data is automatically translated into collection policies, leveraging learned data classifications
- Test against historical data, and measure policy accuracy and coverage before deploying in production
- Review, approve, and deploy instantly

| Policy Name                                                                                                          | Status                                                                                              |
|----------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
|  SSN Blocking Policy                |  Recommendations |
|  Credit card number detection       |  Active          |
|  Prevent credentials and passwords |  Active         |
|  Block source code to USB         |  Active        |



## See and protect data at the prompt.

Get visibility into managed and unsanctioned AI, evaluating prompt content before submission to public models and allowing teams to decide how to address policy violations in real time.

- Discover Shadow AI tools across the organization including users, identities, and subscription type.
- Govern AI interactions with real-time analysis of prompt content, user identity, and intent.
- Prevent leakage, exfiltration and unethical use by alerting on or blocking prompts before submission to AI models.

**"We saw 16,000 events that were false positives, that were then shrunk to a handful of events, which saved me and my team a tremendous amount of time and that was all through Cyera."**

**- Jorge Perez, CISO, Cyera Customer**



# How Teams Use Cyera Omni

## Improve DLP Coverage & Posture

Know exactly how your program is performing - and how to improve it. Identify low-accuracy or noisy policies and gaps in coverage. Omni automatically writes and validates new policies to improve your posture. Keep your tools and maximize your value from them.

## Prevent AI Data Leaks

Stop sensitive data from leaking into unsanctioned or personal AI tools. Omni detects risky prompts and uploads containing sensitive data and can alert, coach, or block in real time. Get visibility across AI tools and reduce exposure without slowing users down.

## Reduce Alert Fatigue

Eliminate the need for manual human review of every alert. AI-driven analysis cuts alert volume by up to 98%, surfacing only the ones that reflect true business risk. Understand the actual sensitivity of data with precise classification, reducing false positives that stem from mislabeled data.

## Detect Insider Threats

Spot early signs of insider data exfiltration by connecting behavioral patterns, access anomalies, and data movement across channels. Distinguish business-as-usual from emerging risk, and surface critical events - even ones other tools classify as low severity.

## Accelerate Investigations

Every alert arrives pre-analyzed: what happened, what data was involved, where it was headed, and why it matters. Get findings summarized in plain language with the business and classification context you need to act. Cut investigation time from hours to minutes.

## Safe Use of Enterprise AI

Enable enterprise AI safely by enforcing company policies inside sanctioned tools. Omni detects regulated or restricted data in prompts and uploads, flags violations, and can block when needed. Keep approved AI available while preventing compliance issues and high-impact data leaks.

## Reduce Third-Party & SaaS Exposure

See what data is being shared externally, who has access, and whether it aligns with approved data sharing policies so you can catch drift and stop risk before it spreads.



# Core Capabilities

## DLP Orchestration

Connect your DLP tools via API. Manage policies and alerts across your stack from one place. Multiple vendors, multiple channels, one view.

## Trends

Recurring activity is grouped into patterns automatically. Spot behavioral spikes, risky data flows, and insider threat signals before they become critical incidents.

## Alert Prioritization

Know which critical and high severity alerts need attention first. See top users, top destinations, and your team's open alert queue.

## Investigation Summary

Get a plain-language summary of every alert: what happened, why it matters, severity, forensic details, and recommended next steps.

## Policy Recommendations

Policies that write and improve themselves. AI learns from your alert patterns, recommends new policies, and refines existing ones. Review, approve, and deploy instantly.

## Policy Analysis & Simulation

Test policies against historical alert data before deploying to production. See projected accuracy and coverage, and validate performance before changes reach your users.

## AI Monitoring & Blocking

Gain continuous visibility into managed and unsanctioned AI tools. Evaluate prompt content before submission and enforce policy in real time based on data sensitivity, identity, and intent.

## Continuous Learning

Every action you take makes the system smarter. Dismiss, adjust, tag, or reclassify. Each input refines how future alerts are scored and prioritized automatically.

## Why Cyera?

High-Speed,  
Agentless  
Discovery at  
Petabyte Scale

AI-Native  
Classification  
with 95%+  
Precision

Data, Identity,  
and Access in  
One Unified  
Platform

Protection for  
Data at Rest, in  
Use, and in  
Motion

Visibility,  
Governance,  
and Defense  
for any AI

For more information visit [cyera.com](https://cyera.com) or book a demo

Book a demo

