GUIDE

Data Loss Prevention Buyer's Guide

Evaluating DLP solutions for the AI era

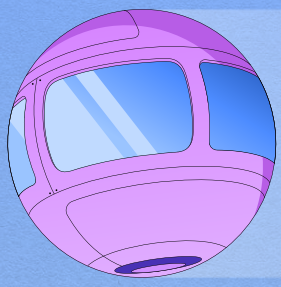

cyera

Table of Contents

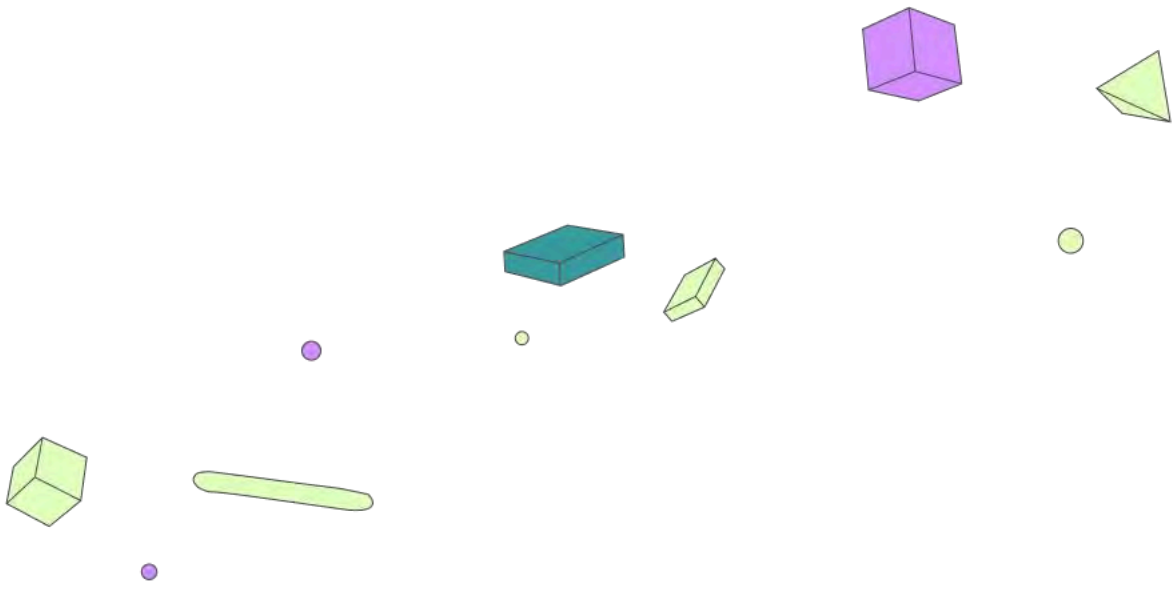

Introduction

You're probably in one of two situations. You already own DLP, maybe several different tools. You have an SSE platform, an email gateway, an endpoint agent, and something native in your productivity suite. Or you don't have DLP yet, and you've been holding off because every peer you've talked to has told you the same thing: buying the tools is the easy part. Operationalizing them is the hard one. Either way, you have two or three analysts, not thirty. Your SIEM queue is already full. Your CISO has asked, in writing, what you're doing about shadow AI. Whatever you decide next has to survive procurement, legal, and probably the audit committee. This guide is written for that situation.

DLP is twenty-five years old. It was built for a world with one email gateway, a web proxy, and users who mostly stayed inside the firewall. That world is gone. Your data now lives across dozens of SaaS apps, multiple clouds, and a growing population of AI tools that read and write data on their own. The architecture underneath classic DLP assumed traffic you could intercept, identities you could pin to a device, and a few chokepoints you could enforce at. None of that holds anymore. SaaS apps move data between each other through APIs that never touch your network. And the enforcement surfaces have multiplied to the point where most teams can't keep policies consistent across them.

The result is the operational reality your peers warned you about. The tools that exist today don't share intelligence with each other. Policies drift between consoles, alerts stack up faster than anyone can triage them, and analysts spend their week closing tickets they already know are noise. When something does go wrong, the evidence trail is thin: you can usually tell that data moved, but rarely who moved it, why, or where it went next. That's the reason teams that already own DLP often don't trust it.

That's the inherited problem. Here's the new one. AI is already on your desk. Agents are starting to access data stores, call APIs, and talk to other agents with no human in the loop. The perimeter is moving from what AI says to what AI does, and your DLP scope has to move with it. That's the question your CISO is asking, and procurement will ask it next.

Better regex won't fix this. Better architecture might. Gartner finds that programs treating DLP as a tool-selection exercise tend to fail. SACR (Software Analyst Cyber Research) goes further: the underlying model has to change, from fragmented point controls to a discovery-led data control plane built for SaaS, cloud, and AI. The fix is intelligence, not another enforcement point. If you already own DLP, you don't have to rip it out. If you're starting fresh, you don't have to inherit the operational burden that turned your peers' programs into ticket factories.

> "The winners won't be the platforms that generate the most alerts. They'll be the ones that measurably reduce risk, improve visibility and control of data with AI and agents, and offer lower operational burden, automated remediation, and audit-grade evidence."
>
> **- Lawrence Pingree, Head of Research, SACR**

What is data loss prevention? A program and technology used to detect, monitor, and protect sensitive information from unauthorized access, accidental exposure, or malicious exfiltration. It governs data at rest, in motion, and in use across endpoints, networks, cloud services, and now AI systems. Modern DLP extends beyond content inspection at fixed enforcement points and is converging with DSPM, SaaS security, browser security, and AI governance.

Architectural Challenges of DLP

Classic DLP rested on architectural assumptions that no longer hold. Data no longer moves through a small number of chokepoints. Sensitive data no longer fits the patterns regex was built to find. Human analysts can no longer keep up with the volume the architecture produces. And AI now acts on data in ways the original model never anticipated. These broken assumptions cascade into five distinct challenges that security teams now operate around. Together, they explain why most DLP programs feel busy without feeling effective.

Challenge #1: Detection that doesn't understand what it sees

Regex evaluates strings, not meaning. A pattern that matches \d{3}-\d{2}-\d{4} flags every nine-digit string with that shape: real SSNs, test values in code, placeholders in documentation, column headers in a schema doc. EDM tightens precision by hashing known values, but only catches data the team thought to fingerprint in advance. Neither approach evaluates the document's meaning, only its surface patterns. The detection engine sees a paragraph about "the Anderson acquisition" and a published press release as the same string and treats them the same way.

An analyst opening Monday's queue is looking at thousands of matches that aren't sensitive at all, and the few that are can land anywhere in the queue, including ranked low severity, because the detection engine has no semantic basis for separating them. Industry analyst Francis Odum reports false-positive rates reaching 90 percent in traditional deployments. Gartner makes a similar observation in its 2025 Market Guide for Data Loss Prevention, noting that traditional content-inspection approaches produce high false-positive rates and performance issues at scale. The architectural ceiling is the same in every case: deterministic detection answers only the questions you knew to ask, at the level of strings rather than meaning.

Challenge #2: Enforcement points that miss the data path

The original architecture was built around three chokepoints: the email gateway, the web proxy, and the managed endpoint. Each assumed the data path was visible to the control: traffic that could be intercepted at the gateway, sessions that could be inspected at the proxy (with all the operational tax of SSL/TLS decryption), and user actions that could be observed by an agent on the endpoint.

Those boundaries don't describe modern data movement. SaaS applications move data between each other through APIs that never touch your network. Cloud data moves between services through API calls and managed pipelines that classic DLP has no visibility into. Users work from unmanaged devices, paste into browser-based AI tools, and collaborate in apps the enterprise didn't sanction. When the enforcement layer doesn't sit on the path, you have no real visibility into the data movement at all, except indirectly through after-the-fact API logs or discovery scans that find the data after it's already moved.

Challenge #3: Policies that can't keep up with behavioral changes

Most DLP policies enforce what you defined in advance: which destinations are allowed, which content patterns trigger action, which user groups have which permissions. When a policy is well-tuned, the tool can tell you who moved what, where, and whether the action matched the rule. What it can't do is recognize behavior that doesn't match a rule but still matters.

A finance analyst emailing a payroll file to the CFO is doing their job. The same file moving from the same analyst to a personal Gmail account at 11 p.m. on a Friday is not. A well-defined DLP policy might catch the personal-domain piece or the after-hours piece individually, but classic DLP can't evaluate the sequence as a behavioral pattern, because it doesn't model what normal looks like for that user. Every event is judged against the rules you wrote, not against the user's own baseline. That's why analysts spend their week closing tickets that fit some rule but represent no real risk, while the events that matter look identical to the ones that don't.

Challenge #4: Tools that don’t share intelligence

Most enterprises run several DLP tools. According to industry analyst Francis Odum, 94 percent of organizations use at least two, averaging more than three.

Consider a departing employee quietly moving customer data out of the company. They export a customer list from Salesforce, save it to their corporate Google Drive, and share it to a personal Gmail account. Salesforce Event Monitoring sees the export. Google Workspace DLP sees the external share. Neither sees the other. An analyst trying to piece the sequence together has to pull data from two consoles and correlate it manually, often days after the data is already gone.

The tools didn't disagree. They just weren't looking at the same thing. Policy intent that was clear in one console drifts in another. Remediation that's automatic in one tool requires a ticket in another. Most teams don't have the operational capacity to keep them aligned, so they don't.

Challenge #5: AI inherits and magnifies existing challenges

AI didn't create new categories of DLP failure. It stress-tested every one of them at machine speed.

The GenAI gap.

Sensitive information now moves through prompts, responses, and AI features embedded inside productivity tools. A developer pasting a public library example into ChatGPT to debug a syntax error is low risk. The same developer pasting proprietary source code with internal logic and credentials is high risk. Classic DLP treats both prompts as the same string of text.

The agentic gap.

The next layer of risk doesn't involve a human at all. Autonomous AI agents access data stores, call APIs, and communicate with other agents at machine speed, with no human in the loop. The primary risk is no longer what AI says. It's what AI does. There is no model in legacy DLP for an actor that isn't a person and isn't tied to a managed device.

What the five challenges share

These look like separate problems. They aren't. Each is a consequence of the same root: DLP was built as a product category, with assumptions about data, users, and enforcement that have all changed. Adding more product to the stack doesn't fix it. The architecture has to change. The next section explains what that architecture looks like, and why most enterprises don't need to rip out what they already own to get there.

What is a DLP Intelligence Layer?

A DLP Intelligence Layer is a centralized decision plane that sits above the DLP enforcement points an organization already operates. It takes fragmented, channel-specific telemetry and turns it into unified, prioritized, evidence-backed decisions, then pushes those decisions back to the underlying tools as policy actions. It doesn't replace existing DLP. It coordinates it.

To do that job, the intelligence layer brings five capabilities together in one place:

- Continuous discovery and classification of sensitive data across cloud, SaaS, and on-prem environments. This is the truth layer that anchors every downstream decision.
- Identity and behavioral context that ties each data interaction to who the user is, what they're entitled to access, how they normally behave, and how data is being shared.
- Risk-based decisioning that ranks events by material business risk rather than rule-match volume, using AI to evaluate content, context, and behavior together.
- Cross-tool orchestration that translates a single policy intent into coordinated actions across the enforcement points an organization already runs, with guardrails and rollback.
- Audit-grade evidence that captures actor, object, action, timestamp, and remediation proof in one investigation-ready timeline.

Cyera's Omni DLP is built on this architecture. So is the broader market shift the major industry analysts have been describing. SACR's Lawrence Pingree calls it the intelligence plane of a modern DLP control architecture. Forrester describes the same direction as the move from "resource-intensive rule-based detection" toward "richer context for confident, risk-based response." Gartner observes that DLP maturity requires centralized intelligence across tools, teams, and data flows.

What this layer is not: it is not a new inspection engine, a new endpoint agent, or a replacement for the DLP tools you already own. It sits above them.

Why Organizations Need One

Three things shift when intelligence is centralized above the enforcement stack.

The alert queue stops being the thing the team works around and starts being the thing the team works from. When the intelligence layer reads every event, including the low-severity ones that never get reviewed, it does two jobs simultaneously. It clears alerts that look dangerous but aren't: a high-severity alert flagging PII sent by an employee to their personal email is a false positive, not an exfiltration event. It also elevates what looks routine: five low-severity alerts about small amounts of PII in individual files can be the signature of a departing employee building a competing company, invisible until the behavioral pattern and the HR record are in the same view. Analysts stop triaging noise. They start investigating risk. Across production deployments, the average alert volume reduction is 80 percent.

DLP policy accuracy is measurable, and most running policies are operating well below what teams assume. A policy running at 20 percent true positive rate is generating false positives or blocking benign activity four times out of five. The intelligence layer measures this, rewrites the policy against historical alerts, and surfaces a replacement tested at 98 percent accuracy before it touches live enforcement. When the team approves it, it deploys directly back into the tools already in the environment. The program stops requiring a specialist to maintain it.

Traditional DLP requires a policy before it can catch a risk. That constraint does not hold when the intelligence layer reads what data is in real time. Source code leaving through an employee's personal account can surface as a risk trend before any rule was written to look for it. The same logic applies to AI tools operating inside the environment. When an AI assistant surfaces information from a confidential document it accessed in the course of answering a routine question, that is not a user error. It is an ungoverned identity with access it should not have. The program can treat AI tools as identities: monitor what they reach, flag access outside appropriate scope, and revoke it. That is the same governance it already applies to human accounts.

Traditional DLP vs. DLP Intelligence Layer

Dimension	Traditional DLP	DLP Intelligence Layer
Architecture	Fixed enforcement chokepoints (email gateway, web proxy, endpoint) built on the assumption that data moves through predictable paths.	Centralized intelligence plane for decisioning and orchestration, paired with distributed enforcement across SaaS, browser, and AI surfaces.
Detection model	Content inspection rules evaluated independently within each channel, with no shared context across email, endpoint, cloud, or SaaS.	AI-driven correlation across data sensitivity, identity, and user behavior, evaluated together rather than channel by channel.
Alert output	High false-positive rates and high alert volume requiring manual triage. Gartner notes content inspection degrades at scale.	Prioritized incidents with data sensitivity, identity context, and plain-language explanations ready for analyst investigation.
Policy lifecycle	Manual tuning on a quarterly or semi-annual cycle, drifts out of date between reviews.	Continuous AI-driven assessment with policy recommendations tied to real alert outcomes, replacing manual review cycles.
Insider risk	Detected at the point of exfiltration in one channel, with no behavioral context to surface earlier patterns.	Cross-channel behavioral baseline surfaces gradual staging patterns before data reaches the last-mile exfiltration event.
AI coverage	Prompt and agent interactions fall outside the enforcement model or are covered by retrofitted controls with limited fidelity.	Natively inspects prompts and responses, discovers sanctioned and unsanctioned AI tools, and governs agent data actions.
Relationship to existing tools	Each tool operates independently, with its own integration, policy management, and investigation workflow. No shared context across enforcement points.	Coordinates enforcement across existing tools and extends coverage to SaaS, browser, and AI surfaces classic DLP cannot reach.

How to Evaluate a DLP Solution for the AI Era

Most organizations already own DLP enforcement points. The question isn't which inspection engine to buy next. It's whether a solution can bring intelligence, consistency, and prioritization to the stack you have.

Each criterion is designed to be verifiable: ask the vendor to demonstrate it, not just confirm it. The criteria are organized into six pillars that map to a modern DLP program.

Capability	What to ask	Why it matters
1. Alert Intelligence		
Alert reduction	Can the vendor show measurable false positive reduction in production deployments, with evidence from live customer environments?	False positive rates in traditional DLP can reach 90 percent - reduction is the single highest-leverage outcome for program credibility.
AI contextual analysis	Does the platform generate plain-language alert explanations with severity, cause, and recommended response?	Analysts shouldn't need to reverse-engineer a rule match. Context-rich alerts cut triage time from hours to minutes.
2. Policy Intelligence		
Policy assessment	Does the platform assess existing DLP policies for accuracy without manual input?	Most teams can't say whether their policies are catching real risk or generating noise. Automated assessment replaces guesswork.
Policy recommendations	Does the platform recommend new policies and refinements based on prior alert outcomes?	Policies drift as the business changes. A learning system closes the loop between alert triage and policy logic.
Policy refinement loop	Does the platform automatically update policy logic based on how alerts are resolved — closing the loop between triage decisions and detection accuracy?	Manual policy reviews happen quarterly at best; automated feedback ensures detection logic improves continuously as the business evolves.
3. Cross-Channel Risk Detection		
Single source of truth	Does the platform provide a unified posture view across DLP channels and providers?	Console sprawl is one of the top operational costs in multi-tool programs - a single truth layer reduces it.
Behavioral baselining	Does the platform establish behavioral baselines per user and role?	Without a baseline, every unusual action looks like every other unusual action. Baselines make prioritization possible.

Capability	What to ask	Why it matters
Data movement trends	Does the platform surface trends in how sensitive data moves across users, channels, and time — not just individual events?	Individual events rarely reveal the full picture; movement trends surface risk patterns before they become incidents.
4. AI Security		
Shadow AI discovery	Does the platform tie AI tool usage to user identity and session context -distinguishing a sanctioned enterprise session from a personal account on the same tool?	The same AI tool carries completely different risk depending on whether an employee is using a corporate subscription or a personal account. Discovery without that context misses the distinction that matters.
Prompt and response analysis	Does the platform analyze AI prompt intent and context in real time - before data is transmitted - or does it rely on pattern matching after the fact?	Regex can't distinguish a productive query from an accidental data leak. Intent-aware analysis at the browser level is the only way to stop exposure before it happens.
5. Usability and Integration		
Identity integration	Does the platform integrate with SSO (Microsoft Entra) and identity context for user-level risk scoring?	Identity is the connective tissue between data, behavior, and risk. Without it, every alert is anonymous.
DSPM integration	Does the platform integrate with DSPM to identify sensitive data at rest so monitoring and prevention policies are in place before that data starts moving?	DSPM gives you a head start. Knowing where sensitive data lives before it moves means your policies are already in place when it does.
Coexistence with existing tools	Can the platform coexist with current DLP providers during a phased rollout?	Rip-and-replace is rarely feasible. Coexistence lets teams consolidate intelligence while preserving enforcement investment.
Time to value	Can the vendor demonstrate initial visibility in 30 minutes and actionable policy insights within 14 days?	AI threats and insider risk are active now. Quarter-long integrations leave organizations exposed during the critical window.

Getting Started Without Starting Over

Starting does not require ripping out existing tools, completing a data governance program first, or having a mature classification foundation in place. The intelligence layer integrates with what is already running and builds the intelligence layer above it. Programs at different stages of maturity enter at the same starting point: connect the tools you have, let the engine assess what it finds, and go from there.

The program builds in four stages. Most organizations start with what their users already rely on every day: the productivity and collaboration environments where most data actually moves. Discovery catalogues what sensitive data exists in those environments before enforcement begins. For organizations without an existing classification standard, the program provides one as a foundation rather than requiring one as a prerequisite. Governance follows, establishing policies and testing them against historical events before anything goes live.

Real-time enforcement is the third stage, applied across channels in proportion to risk. The fourth stage is the one that changes the conversation with the business. With a functioning intelligence layer in place, security shifts from blocking AI adoption to governing it. Teams move from saying no to new tools to showing how those tools can be used safely.

What used to take years to build now takes quarters, because the work that consumed the most time is automated.

"Omni DLP has really been able to, through AI, identify what's really important, what looks like normal business processes versus things that seem abnormal. It has helped us focus on things that matter. We saw 16,000 events that were false positives, then shrunk to a handful of events, which saved me and my team a tremendous amount of time. That was all through Cyera."

- Jorge Perez, CISO, Financial Institution

Frequently Asked Questions

Does an intelligence layer replace DLP tools?

No. The intelligence layer sits above the enforcement points you already operate and orchestrates them. It reads telemetry from existing tools, enriches it with data and identity context, and pushes policy decisions back to those tools. Nothing needs to be ripped out.

DLP is already built into my productivity suite. Why do I need this?

DLP controls in productivity suites like Microsoft 365 or Google Workspace cover the application ecosystem of that vendor, including email, endpoint, cloud drives, and collaboration apps, but only within it. What they don't see by design is the data that crosses that boundary: insider threat spanning other channels, behavioral signals from enforcement points outside the ecosystem, and AI data paths that route beyond the vendor's perimeter. The intelligence layer reads across all those enforcement points and surfaces the risk picture no single vendor's tools can assemble.

Why does DLP produce so many false positives?

Because classic DLP matches patterns, not meaning. Regex checks string shapes; fingerprinting catches values you remembered to hash. A file flagged because it contains a 16-digit number gets the same alert whether it's a credit card or a product ID. AI-native DLP is semantic: it evaluates what a document means, who is interacting with it, and how that person normally behaves. That shift from pattern to context is what brings false-positive rates down from the 90 percent range traditional deployments produce to something a small team can actually work from.

What's the difference between DLP and DSPM?

DSPM or Data Security Posture Management finds and classifies data at rest. DLP enforces what happens when anyone interacts with it (i.e., data in motion and in use). Run one without the other and there's a gap: DSPM without DLP gives you visibility but no enforcement; DLP without DSPM gives you enforcement against patterns rather than known sensitive data. If you already own DSPM, that classification work carries directly into DLP deployment. If you don't, the intelligence layer builds the classification foundation as part of onboarding rather than requiring it as a prerequisite.

How does this relate to my insider risk program?

Insider risk programs and DLP address overlapping problems from different angles. Insider risk focuses on user behavior and intent over time; DLP focuses on what happens to data at the point it moves. The intelligence layer connects both. It builds behavioral baselines per user and role, surfacing gradual staging patterns before they reach the last-mile exfiltration event. If you already run an insider risk program, the intelligence layer gives analysts richer, cross-channel context to work from rather than running as a separate silo.

Already own DLP? Make it work together. No DLP? Get it right from the start. See how with Cyera Omni Data Loss Prevention.

Get a Tour →

About Cyera

Cyera is the AI Security Platform built for the age of agents. Enterprises like Paramount, Chipotle, and Valvoline use Cyera to control exactly what data their AI can reach — and govern what happens next. The platform secures data at rest, in motion, and in use, whether touched by humans or AI agents. Valued at $9 billion and backed by over $1.7 billion from Accel, Blackstone, Cyberstarts, Georgian, Lightspeed, and Sequoia.

Learn more at www.cyera.com, or follow Cyera on LinkedIn.

cyera