SACR®

APRIL 2026

# THE GREAT DLP RESET

# SECURING DATA IN THE AGE OF SAAS, CLOUD, AND AI

LAWRENCE PINGREE

ORION

# We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

## About Software Analyst Cybersecurity Research

SACR is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over 80,000 readers and followers, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

### Author


- Lawrence Pingree is the Head of Data and AI Security at SACR, where he leads research on data protection, AI security, and agentic security models. He brings more than ten years of analyst experience from Gartner and has authored over 300 research notes across cloud security, endpoint defense, SD-WAN, and AI security.

# Table of Contents

# Executive Summary

## Data Loss Prevention (DLP) is undergoing a structural reset.

What was once a fragmented set of point controls anchored to endpoints, networks, and email gateways is being rebuilt into a unified, discovery-led data control plane designed for a world where data no longer sits still, users no longer operate within defined perimeters, and AI systems now participate in how data is created, transformed, and exfiltrated.

The legacy DLP model is breaking, not because detection has failed, but because the underlying assumptions no longer hold. Traditional DLP relied on stable data patterns, centralized enforcement points, and human-driven tuning. Modern environments defined by SaaS sprawl, cloud data gravity, and GenAI workflows have invalidated each of these assumptions.

Our report argues that the next era of DLP will not be defined by better pattern matching or broader coverage, but by a fundamental architectural shift.

1. **Our thesis for this research is that Data Loss Prevention (DLP)** is being rebuilt into a discovery-led control plane for modern runtimes (SaaS, cloud data, and AI/agent workflows). Traditional Enterprise DLP, Integrated DLP, and Cloud-Native DLP categories are an insufficient way of delineating DLP at the program level. A DLP reset is needed. In the DLP reset, the winners are not the platforms that generate the most alerts, but those that reduce operational burden while delivering measurable risk reduction through automated remediation, real-time prevention where feasible, and audit-grade evidence (and lineage where possible).
2. **DLP is being rebuilt into a discovery-led control plane:** This includes continuous classification, which becomes the truth layer where identity/entitlements provide decision context, and enforcement shifts from a few chokepoints to a set of distributed enforcement planes (SaaS APIs, inline SSE/SASE/GW, browser/session, endpoint, and AI prompt/agent surfaces).
3. **The Winners:** The winners won't be the platforms that generate the most alerts. Rather, 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.

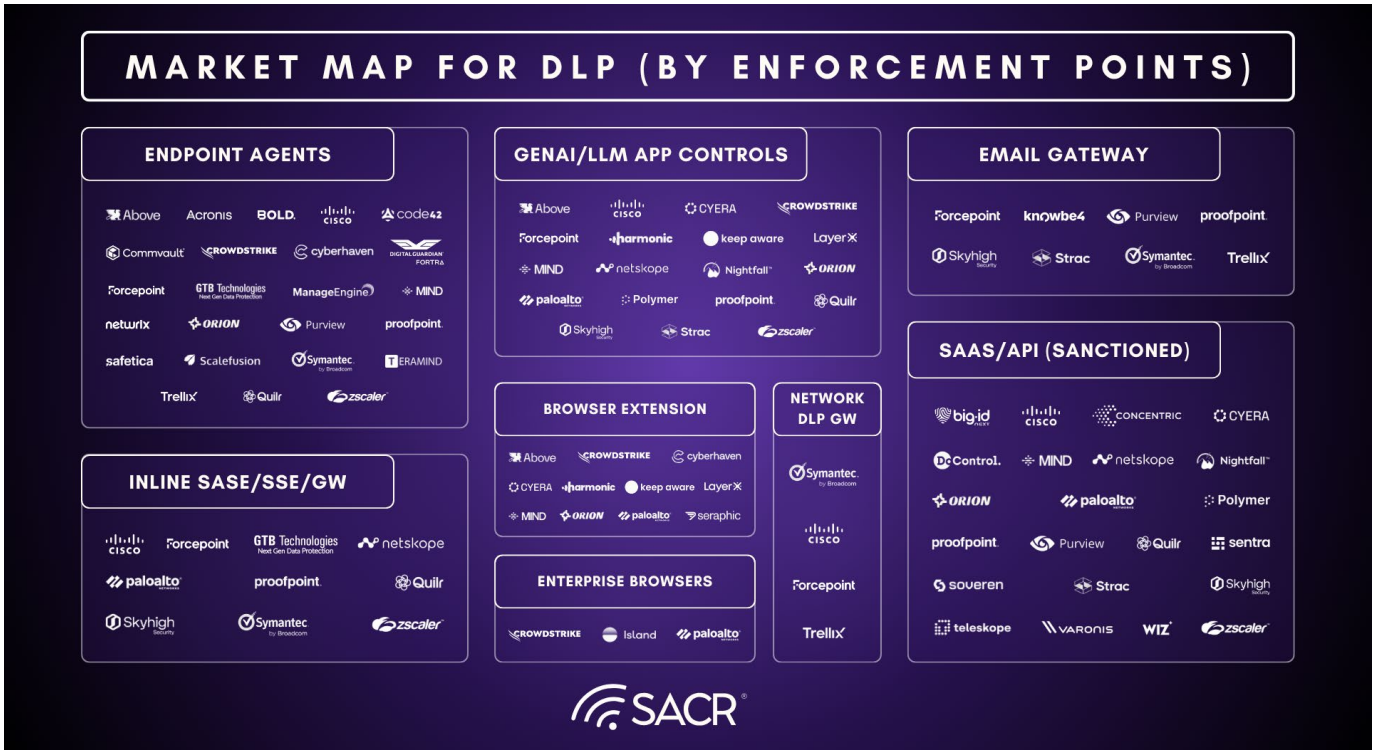

# Key Insights On How DLP Is Changing

DLP is undergoing a reset and moving from deterministic, enforcement-point-centric controls to a discovery-led data control plane model that combines continuous classification, identity, context, data labeling, and automated remediation across SaaS, cloud data, endpoints, browsers, user-driven and agentic AI workflows. AI adoption has reached 73% of enterprises in 2026, while real-time security governance is just beginning to emerge at 7%. Browsers are emerging as a key defense mechanism for AI and Data loss as users spend ~75% of their work day either working in a web browser or attending virtual meetings.

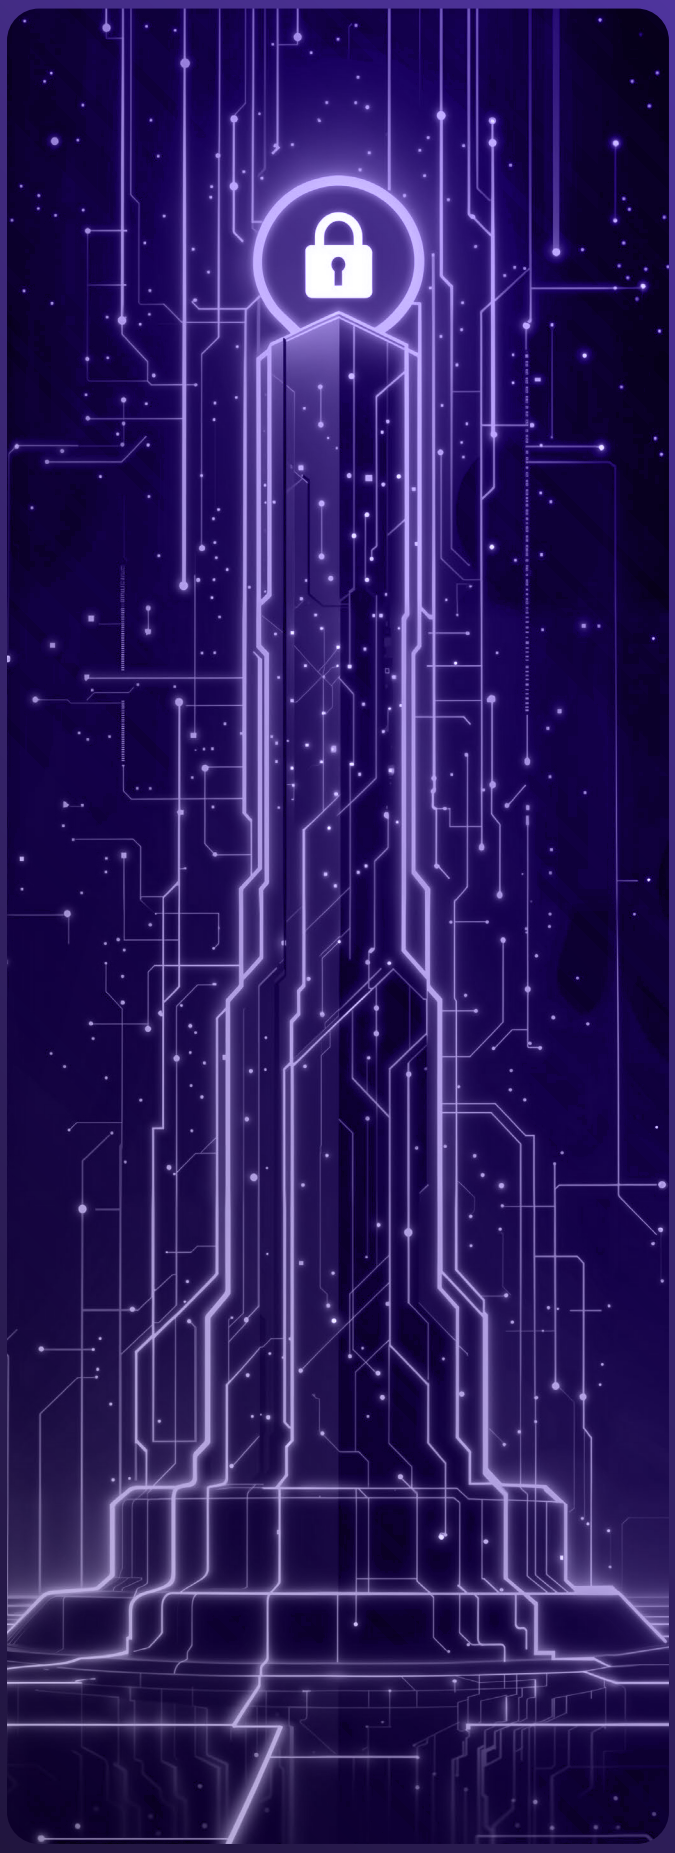

Based on discussions with SACR clients, vendor briefings, and public information:

- The center of gravity is shifting toward faster time-to-first-signal (TTFS), faster time to prevention and much more focused on lower tuning burden to avoid false positives, and improve classification.
- API, SaaS-first and DSPM-led control planes generally offer lower burden, while classic endpoints and network suites remain high-burden even when their breadth is strong.
- Inline SSE and SASE DLP (e.g., Prisma Access, Netskope, Zscaler, etc) remains a high-value path for broad enforcement, but it carries structural burden, for example requiring traffic steering or custom reverse proxying, policy pattern and profile management, and dependency on traffic flowing through enforcement points.
- M365-native DLP (for example Microsoft Purview) can deliver high value quickly inside Microsoft environments, but the operational tax includes activities such as endpoint onboarding, policy complexity, and false positive or custom data classification tuning which remains non-trivial.
- Expanding Agentic Platforms and AI interaction and enforcement points is complicating the future of DLP especially for emerging agentic platforms and their workflows.

# Actionable Summary

- **Establish a truth layer for the data security program (discovery/classification):** If you cannot answer where sensitive data is and who can access it, enforcement will be noisy and politically hard to sustain and preventive remediation efforts will not be easily achievable.
- **Add in essential context (identity, entitlements, sharing posture):** Combine a truth layer (DSPM and discovery) with context (identity and entitlements) to inform enforcement and direct remediation efforts or automations (for example data owner guidance or automated approvals).
- **Choose enforcement points intentionally (SaaS API first where possible; inline and endpoints where justified):** Utilize a targeted enforcement point layer (for example, by using SaaS APIs, Inline SSE/SASE/GW where needed, and focus on endpoints only where necessary - e.g. users leveraging traditional protocols and data sharing infrastructures like SMB/SAMBA Drive Shares or sync-share file solutions).
- **Operationalize GenAI runtime controls, ideally using browser controls (prompt/session + agent tool-call logging):** Define controls for users for example prompt pasting, Copilot access scope, and agent tool-call auditing and logging and don't assume classic DLP policies will translate to the new emerging areas.
- **Measure data loss prevention outcomes (automation, risk reduction and adequate evidence):** The reset winners are those that can revoke sharing, redact, delete or even mask and encrypt content in SaaS, and produce evidence trails with less analyst labour and without increasing burden or creating a ticket factory outcome for your human participants.
- **Focus on Audit-grade evidence and end-to-end data lineage (what counts):** The key components are the event logs, the actors involved, the object impacted, the action taken, the precise timestamp, data lineage trace and proof of the remediation if applied.

## Introduction, Market and Industry Context

# The Great DLP Reset

The Great DLP Reset signifies a major transformation in architecture, detailing several essential shifts in how data security initiatives must evaluate DLP tools. This movement is steering the development of new vendor strategies to resolve longstanding challenges within data loss prevention programs. DLP is transitioning from conventional, static perimeter security toward a discovery-driven data control plane, tailored for the rapid pace of Cloud, SaaS, AI, and autonomous agent interactions.

## Market Definition: Modern DLP

Modern DLP (as used in this report) is defined as solutions that, when combined, deliver a set of capabilities that:

Data Loss Prevention (DLP) is a security strategy and set of technologies designed to detect, monitor, and safeguard sensitive information from unauthorized access, accidental exposure, or malicious exfiltration. It governs sensitive data at rest, in motion, and in use across endpoints, networks, and cloud services. This definition intentionally expands beyond content inspection at fixed enforcement points and reflects the market's convergence with DSPM, SaaS security and governance, browser security, and AI governance.

### Definitional Technology, Feature(s), and Service Lines

DLP tools perform automated data discovery and classification with deep content inspection, and contextual policy enforcement (block, quarantine, encrypt) focused on the enforcement of real-time data prevention on real-time user interactions. Integrations commonly include CASB/SSE/SASE, network gateways and endpoint agents to govern data flows. Emerging patterns include browser extensions and agentic platforms with limited data control and visibility. Modern DLP tools prevent sensitive data exfiltration across endpoints, SaaS, web/email, and cloud via classification and policy enforcement.

***Exclusion criteria: Vendors that deliver products and services of AI prompt inspection and API based application proxies*** *that focus primarily on prompt threat inspection and context (these solutions are more focused on agentic and workflows: See Unified Agentic Defence Platforms (UADP).*

# What is This Reset in DLP?

**The architecture of the DLP reset defines a new truth layer combined with context and distributed enforcement, establishing continuous discovery and classification to determine what the data is and where it is, with enforcement that includes remediation and substantive evidence.**

- **Truth Layer:** Extends across premises, SaaS, cloud, AI services and endpoints.
- **Context:** Integrates identity, entitlements, sharing posture, and behavioral analytics to provide decision intelligence and reduce false positives.
- **Distributed Enforcement:** Shifts control from fixed chokepoints to multiple planes including SaaS APIs, inline SSE/SASE/GW, browser sessions, and AI prompt surfaces.
- **Remediation:** Emphasizes automated actions such as revoking links, redacting sensitive data fragments, engaging data owners through autonomous workflows and chat applications (Slack/Teams) and can perform quarantining of assets at machine speed.
- **Evidence:** Produces audit-grade logs and data lineage traces that provide proof of remediation and a forensic chain of custody for investigations in the case of an intentional data exfiltration by a user.

# What It Isn't

Rather than a traditional solution defined by pervasive regex, the DLP reset moves away from models that necessitate constant manual adjustment and result in a ticket factory style SOC environment. Such legacy approaches impose a heavy operational burden on personnel at every stage, from policy refinement and alert triage to the execution of enforcement actions.

- **Regex Everywhere:** Traditional DLP relies on deterministic pattern matching and static signatures (e.g., RegEx for Social Security Numbers) which fail to understand content depth or context in modern, unstructured data flows.
- **Manual Tuning:** Legacy systems require constant human intervention to manage exception sprawl and adjust rules for stable data schemas that no longer exist in ephemeral cloud workloads.
- **Ticket Factory:** Without identity or behavioral context, classic tools generate massive volumes of false positives, turning security operations into ticket factories focused on bulk-closing alerts rather than active risk mitigation.

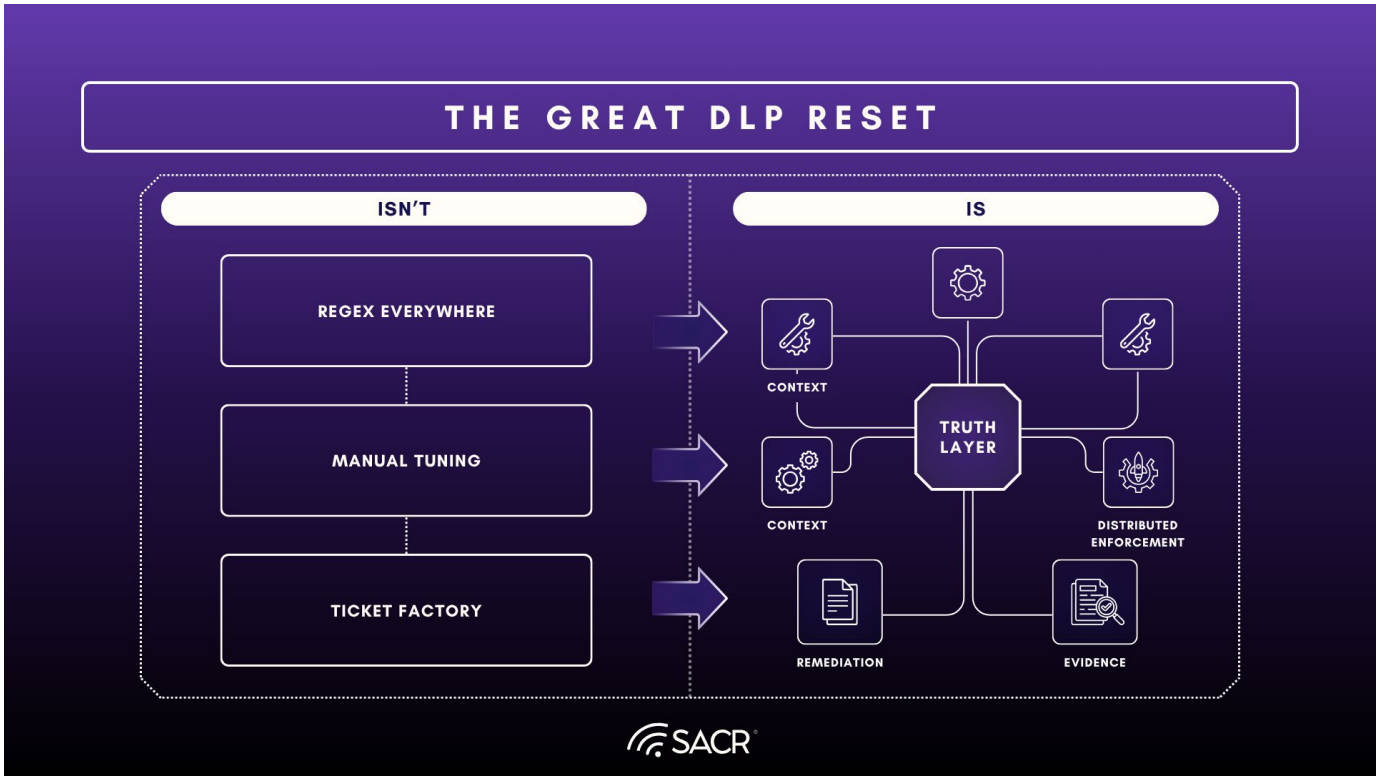

## The DLP Control Plane (Discovery and labelling-led)

Modern DLP is best understood as a control-plane model which includes decisioning and orchestration that drives actions across multiple enforcement planes. It breaks the regex everywhere plus manual tuning era and ticket factory pattern by separating the intelligence plane that determines what matters, from the enforcement planes that execute data security control. Put simply, data discovery, a core aspect of DSPM technology, is great at discovery and classification of data, determining location and various use cases. DSPM tools and the data they discover and classify help inform enforcement points and the placement of key inspection and control layer functions across an enterprise, its users, and any AI agents it is operating. It can be used to fine tune data loss programs for actual data loss prevention, vs only being used like they are in most DLP programs, as data monitoring solutions, or focused on limited data classifications such as PII or well known structured data types.

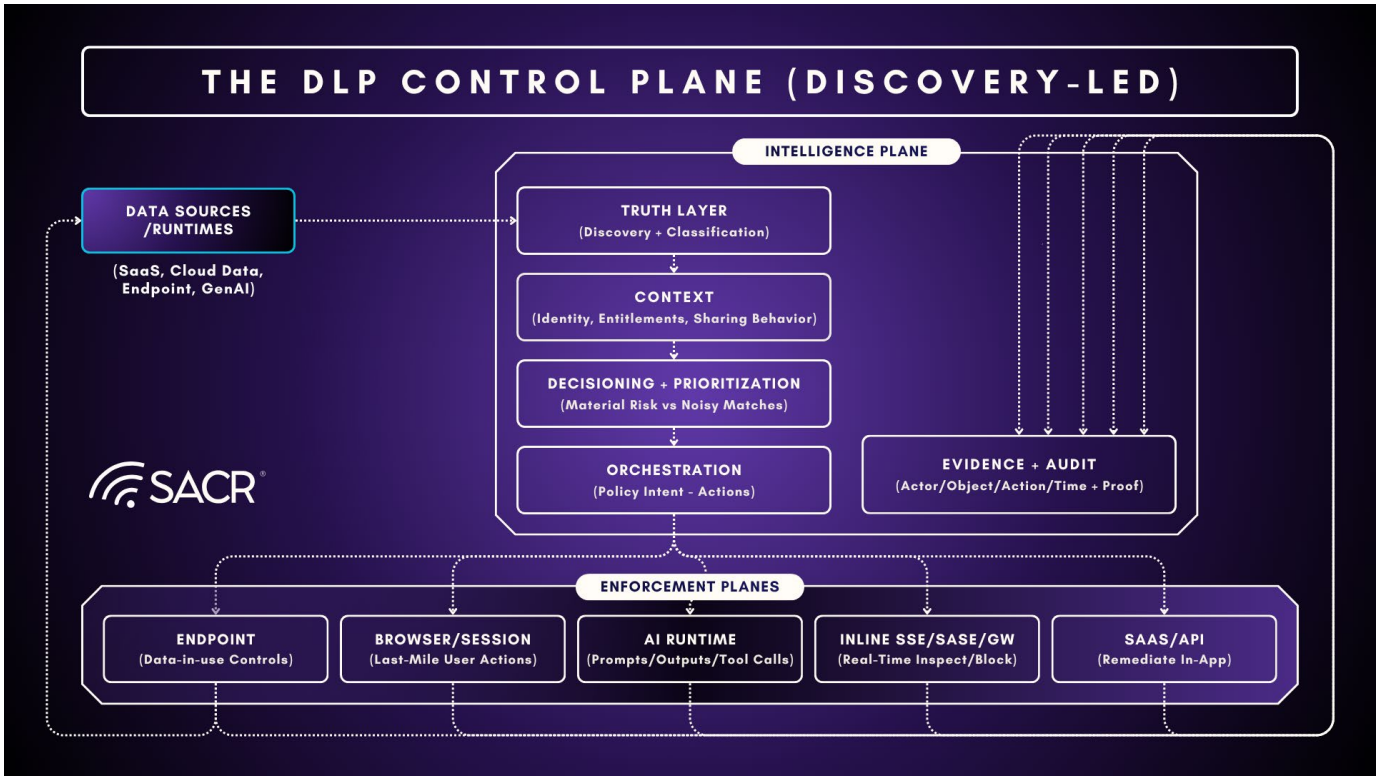

# The Intelligence Plane (truth, context and orchestration)

Modern DLP Control plane architecture is fundamentally sub-divided into two major components: **the Intelligence Plane and the Enforcement Plane**. This structural division acts as the core machinery of the reset, effectively separating the centralized intelligence required to determine data significance from the distributed planes required to execute precise controls.

- **Truth layer (continuous discovery and classification):** Establishes what the data is, where it lives, and how it is moving. Where vendors offering Data Security Posture Management (DSPM), handle data discovery across premise and cloud, properly classify and label, and synchronize data sensitivity labeling across Microsoft (Purview/MIP) and Google Workspaces. These labels serve as critical truth layer context for the enforcement planes.
- **Context layer (identity, entitlements, sharing posture and behavior):** Determine who can access it, how it's exposed, and what the risk scenario is to apply the best controls to match the scenario.
- **Decisioning and prioritization:** Reduces noise by ranking what is materially risky (not merely what matches a simple pattern), augmented with AI and LLM content analysis and labeling.
- **Automation and orchestration:** Translates policy intent into repeatable orchestrated actions (for example warn, block, redact, revoke sharing, quarantine, label, encrypt) ideally with guardrails, contextual policy control and rollback capabilities.
- **Evidence and auditability:** Produces investigation-ready data lineage and timelines with actor, object, action, timestamp, and remediation proof (plus end to end data lineage where possible).

# The Enforcement Plane (aka distributed control points)

The Enforcement Plane represents the distributed control points where security policies are actively applied to data across various environments. These planes are responsible for executing specific actions, such as blocking, redacting, or quarantining, at the precise moment a violation occurs, moving beyond simple detection to active data risk mitigation. By shifting enforcement from a few centralized chokepoints to distributed surfaces like SaaS services, APIs, inline SSE/SASE/GW, and the enterprise browser, organizations can achieve more precise control over modern data movement and usage patterns. SACR believes that enterprises must look at their data security architecture through this lens to achieve better data loss prevention outcomes.

The **Enforcement Plane** includes:

- **SaaS/API enforcement:** Out-of-band controls and remediation inside collaboration and business apps and in AI workflows and agentic platforms.
- **Inline SSE/SASE/GW enforcement:** Real-time inspection/control for web and SaaS traffic where steering and SSL/TLS decryption is justified.
- **Browser and session enforcement (last-mile runtime):** In-session controls over the dominant leakage verbs (copy/paste, upload/download, printing, screen capture), including browser based GenAI interactions.
- **Endpoint enforcement:** Device-level controls for local exfil paths (USB, local copies, unmanaged sync, print, clipboard, screen capture), especially for regulated or high-risk endpoints and end user workspaces.

## Emerging Problems in Data Security and Data Loss Prevention

Classic DLP approaches struggle because they assume clear architectural or isolated choke-points, stable data schemas and patterns, and that there are human resources available for manageable tuning and remediation. Modern environments violate those assumptions dramatically, leading to failed deployments and fragmented enforcement with no clear unification of policy.

Organizations are facing a modern data exposure problem characterized by:

- SaaS sprawl and collaboration-first workflows for example sensitive data moves through Slack, Microsoft Teams, Google Drive, One Drive, Microsoft GitHub or various SaaS applications like Salesforce, marketplaces and similar SaaS surfaces.
- Cloud data gravity has shifted towards sensitive data in warehouses, data lakes and object stores with complex access paths and various data sharing integrations (even between SaaS applications).
- User controlled GenAI chat applications and agentic workflows, prompt-based exposures, copilots with broad reach and data moving data between data stores, file systems, applications or other AI agents and users.
- Expanding SaaS-based agentic platform services and no-code providers (Salesforce Agents, Claude Cowork, OpenClaw hosting, Eigent(open source), Zapier, Airtable etc)

# Why Classic DLP Became High-Burden

Classic Data Loss Prevention (DLP) approaches have become a high-burden due to several core problems that fundamentally require a reset. These issues include a significant operational tax from noise and tuning burden caused by high false-positive rates and exception and enforcement gap sprawl. Classic DLP suffers from weak visibility into where sensitive data exists and which users have the entitlement(s) to reach it. Often various DLP tools were deployed in silos and not properly configured to unify enforcement or consistent policies and utilized various editions of data labeling and classification techniques. Architecturally, it has a poor fit for modern environments characterized by SaaS-native sharing and API-driven data movement that bypass intermediary choke points. Classic DLP is proving to be even a weaker fit for AI usage, struggling to govern probabilistic AI workflows, prompt-based exposures, and agentic tool calls.

While these foundational mechanics still have a place for performance-heavy compliance tasks, they are insufficient for modern runtimes and emerging architectures which are requiring greater granularity and data-in-use understanding. This legacy framework was effective only as long as the enterprise controlled the infrastructure, the network, application deployment architectures and the devices. However, the shift toward Software-as-a-Service (SaaS), Cloud-native infrastructures, and now Artificial Intelligence (AI) has fundamentally disrupted this paradigm.

The Core Problems and Operational Tax

- **Noise and Tuning Burden:** High false-positive rates and exception sprawl leading to DLP-as-a-ticket-factory style outcomes, overwhelming data security teams and causing alert fatigue for security operations.
- **Weak Visibility and Graph Visualizations for Analysis:** Traditional tools lacked insight into what sensitive data exists at rest and who has the entitlement to reach it across fragmented environments. They lacked significant visualizations and graph databases to properly articulate vast interconnected data flows between various entities and connect it with data in-use visualizations.
- **Architectural Fragility:** Dependence on network interception makes steering traffic and rearchitecting for DLP fragile and politically difficult as data moved directly between SaaS apps via APIs (for example N8N and Zapier style no-code and workflow automation tooling).
- **Weak AI Fit:** Static rules cannot govern emerging probabilistic AI workflows, prompt-based exposures, or agentic tool calls and need additional enhancements to deliver that functionality through integrated APIs.

# Market Shift Timeline:
# Data Loss Prevention Challenges 2000s-Present

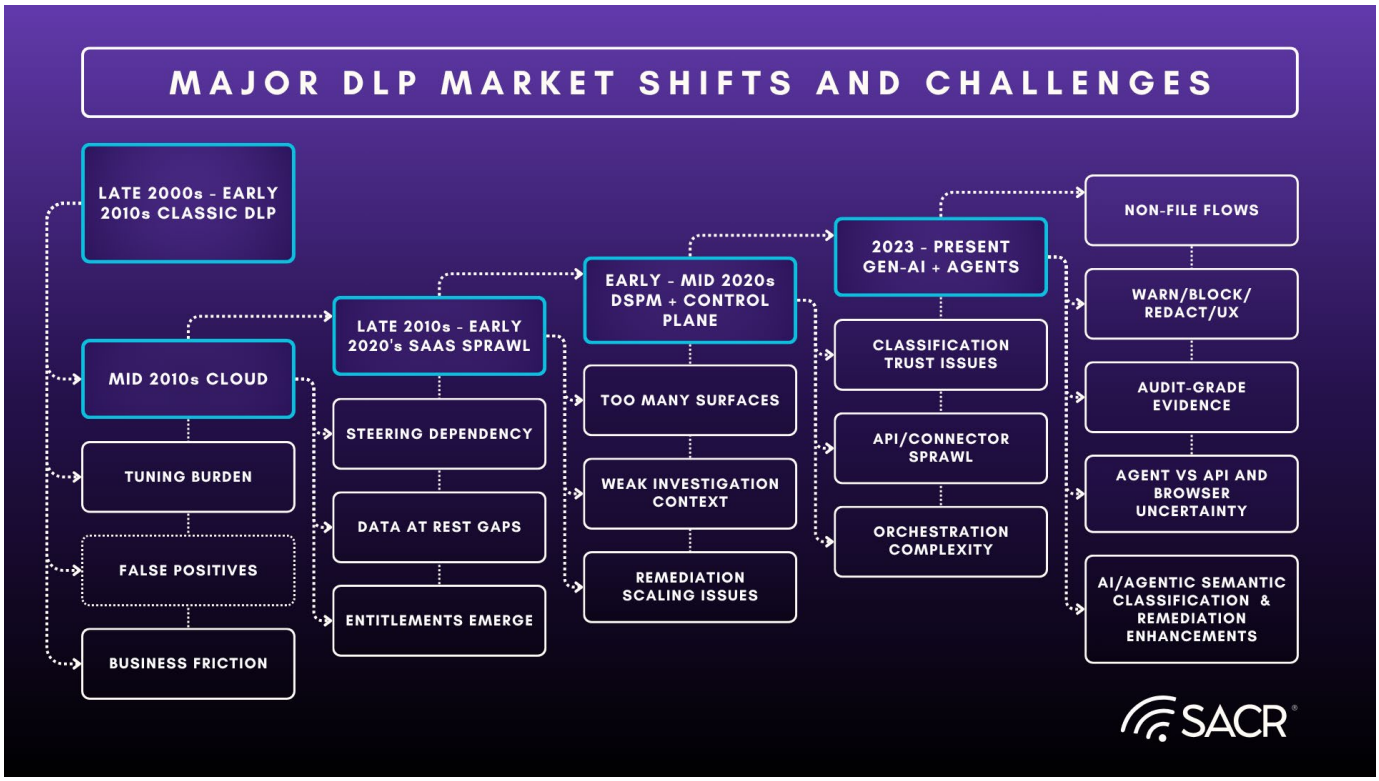


## Late 2000s to early 2010s: DLP becomes a Mainstream Enterprise Control

- **High policy authoring and tuning burden** (regex, Exact data matching/hashing (EDM), fingerprinting) and large operational overhead
- **High false positives without strong context**, leading to alert fatigue and DLP-as-a-ticket-factory
- **Cultural and adoption friction:** DLP often perceived as blocking business productivity without clear outcomes

## Mid 2010s: Cloud migration begins to erode chokepoints

- **Architectural dependency** on interception of network traffic, steering became fragile and politically difficult.
- **Visibility gaps for data-at-rest across hybrid** environments and in SaaS and cloud storage (not just data-in-motion).
- **Identity use and entitlements** started to dominate outcomes, but classic DLP has weak native entitlement context and is mostly limited in scope to traditional endpoint or network traffic choke-points, file and data storage systems and sharing protocols.

## Late 2010s to early 2020s: SaaS sprawl and collaboration-first work explode the policy burden

- **Proliferation of enforcement** surfaces (too many places to write and maintain consistent policies)
- **Weak evidence and investigation context** (who shared what, with whom, and why)
- **Remediation becomes workflow-heavy** (revoking shares, cleaning repos, fixing permissions) and doesn't scale well via manual tickets

## Early to mid 2020s: DSPM and discovery-led approaches reshape expectations

- **Classification Accuracy limitations:** Classification accuracy and explainability at scale (trust becomes the gating factor) for adoption to be viable.
- **Shadow SaaS (also called Shadow IT) and SaaS data Sprawl:** Connector and data sprawl and integration reality (coverage of data in storage and motion becomes dependent on APIs, audit logs, and action depth) often requiring knowledge of SaaS adoption to integrate DLP functions.
- **Control-plane complexity:** Orchestrating consistent actions across heterogeneous tools without breaking business workflows meant that end to end visibility and discovery challenges became rampant.
- **Centers of gravity emerge for data classification labeling:** Microsoft Information Protection/Purview, Google Workspace become centers of gravity for key data labeling hierarchies, enabling greater federation of enforcement across the enforcement layers

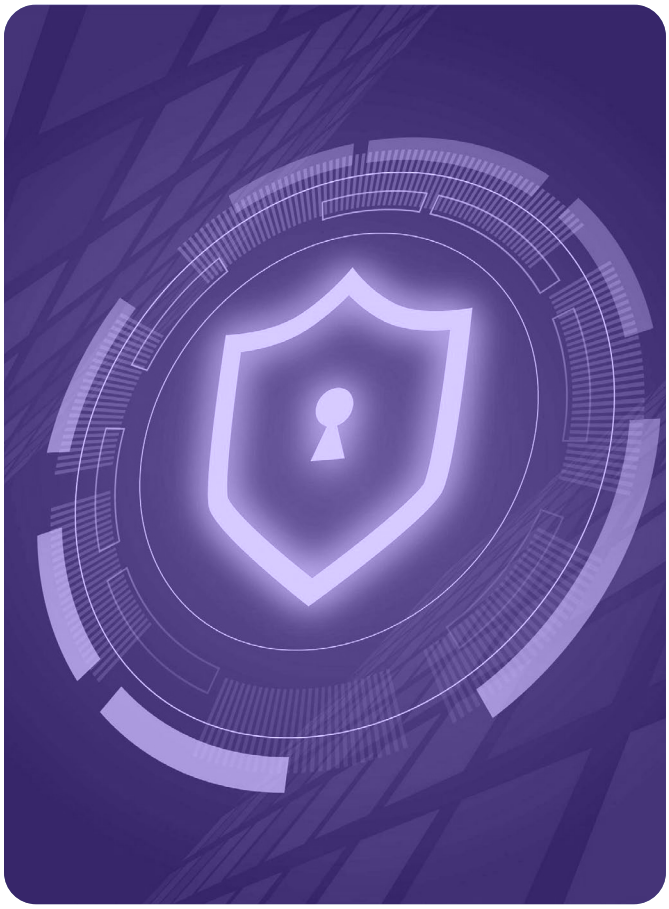

## 2023 to present: GenAI and agentic workflows create a new DLP runtime

- **Governance:** Governance of non-file data flows (prompts, outputs, tool calls) where content inspection at gateways is insufficient.
- **New UX and policy questions:** Warn user vs block or data redaction require deep user and use case understanding, and what constitutes sensitive in a prompt context is required for proper enforcement and policies.
- **Evidence and audit requirements:** Logging agent actions and maintaining chain-of-custody for investigations and having a forensically sound chain of custody and data lineage become a new requirement.
- **Model and AI agent uncertainty:** Controlling stochastic systems and verifying efficacy claims in production become very challenging and introduce uncertainty.

# The Great Reset: Why DLP is Back

We think DLP programs need to reframe their deployment architectures around our new framework. Today, DLP deployments span a variety of enforcement plane capabilities and fragmented abilities to control data (also called actions), some using older or more basic methods of data identification, classification and enforcement. Meanwhile, more modern classification tooling such as that of machine learning, semantic classifiers and the broader variety of data security actions can be embedded across enforcement points in a more unified manner. The old frameworks were moderately effective because the enterprise controlled the infrastructure, the network, the communication channels, and the computing devices that made data residency and transit patterns static and manageable.

Though the basics of traditional DLP still have a place in modern DLP systems (for example using Exact data match or regular expressions, for delivering higher performance or speed), they are not the capabilities that offer the best detection accuracy, classification depth and deep semantic understanding of content and context that is needed for future data loss prevention, especially agentics. Today's disparately deployed DLP solutions also need a better grounding in truth, a new defining truth-layer to align various enforcement points and their policies. Modern DLP implementations are plagued by severe architectural instability and have various nuanced limitations. Security departments are frequently trapped in a cycle of managing inconsistent, siloed controls that span across endpoints, network gateways, collaboration platforms, and browser environments. This disconnected approach necessitates a comprehensive structural overhaul, shifting toward a unified control plane anchored by a discovery-driven data strategy. Failure to implement this transition prevents organizations from maintaining uniform policy enforcement across the decentralized points needed to regulate real-time data interactions by modern users and emerging AI agents.

## Emergence of the Ticket Factory Overload

In traditional DLP solutions, this phenomenon is usually referred to as alert fatigue, or false positive overload, or as the swivel-chair security problem - where practitioners need to stitch various events and data security contexts across multiple silo's of tools. But generally, this is why traditional DLP tends to turn security teams into ticket factories, and in the new world of DLP the core idea is to avoid it. Traditional DLP solutions rely on static rules, regular expressions (regex), and exact data matching (EDM). For example, if a user tries to move a file under the old classification schemes where it looks like it has 16 digits, the regex based DLP flags it as a credit card number, even if it's not, simply because the data is formatted similarly. This is predominately because these older tools lacked context and depth of understanding of semantics and consideration of other data content context, they often generate massive volumes of false positives. For example, a 2024 Security Boulevard SOC Efficiency Study noted that nearly one-third of all security alerts are false positives, and legacy DLP classification approaches were a primary offender.

## The Old Paradigm Created a Tremendous Toll on the SOC

In the old ticket factory paradigm, analysts with DLP tooling and enforcement points without solid context often spend their entire day bulk-closing tickets just to keep their heads above water, investigating business-as-usual activities rather than actual data exfiltration events. The SOC or data security teams become a ticket factory focused on closing IT service desk requests rather than hunting threats. As data classification, context and semantic understanding has increased, so too has accuracy of alerting, and a great reduction in false positives.

## Challenges in Modern Day Environments

- **SaaS Sprawl:** Sensitive data now moves from both managed and unmanaged endpoints, through communications tools like Slack, Microsoft Teams, and GitHub and are also often outside the view of traditional DLP style gateways and choke points.
- **Unified Data Labeling Strategy and Synchronizing Across the stack:** Google data labeling or Microsoft Purview, tend to be gold standards. Using modern DSPM tools like Cyera, Palo Alto Networks, Varonis and others, they can properly translate labeling as files or data move between different environments.
- **Cloud Data Gravity:** The emergence of Data lakes and object stores has created new complex access paths that violate the assumptions of legacy choke-point models.
- **The AI runtime increasingly routes through the browser:** Most SaaS work and many GenAI interactions happen in-browser, which makes browser and session control a first-class last-mile enforcement-plane. Endpoint and network DLP are not obsolete, but they are often insufficient alone for browser-mediated leakage (copy/paste, upload/download, printing, screen capture) and for prompt and agent workflows.

# Trends driving the change

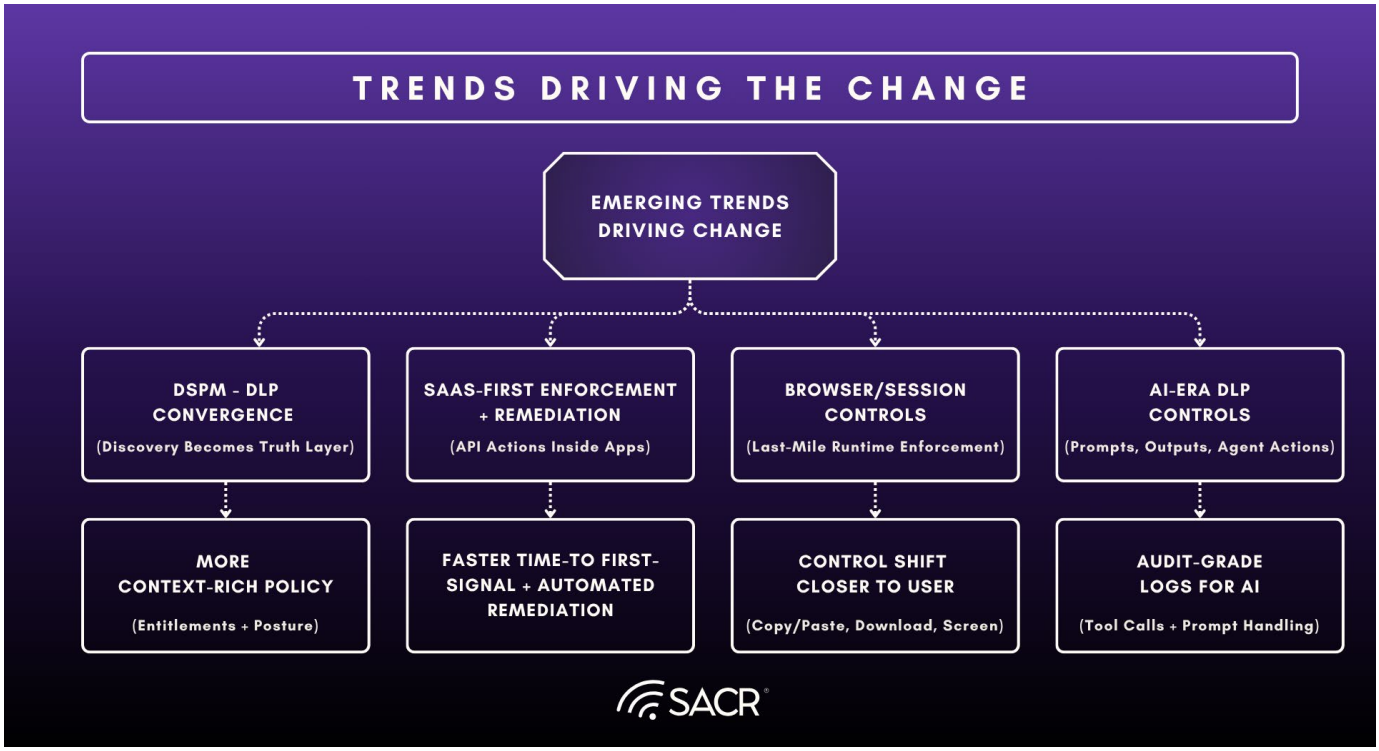


These shifts are necessary to address modern environmental challenges, including **SaaS Sprawl**, where sensitive data moves through tools like Slack and GitHub outside traditional gateways, and **Cloud Data Gravity**, which involves complex access paths in data lakes and object stores and labeling problems. As the **AI runtime** increasingly routes through the browser, traditional endpoint and network DLP are often insufficient to control browser-mediated user actions like copy/paste and AI prompt interactions. It's notable that not all enforcement points have orchestrated data control, nor do they all share classification and labeling, an industry standardization problem that some vendors are addressing. To address this, vendors focused on data loss have been centralizing policy enforcement based on document and file labeling within data classification functions (for example, labels provided by Microsoft Purview or Google Workspace) in the enforcement layer, most providers now offer label based enforcement, helping to unify enforcement actions across the enforcement layer. To address the data relabeling issue, some enforcement providers offer relabeling capability for documents that pass through their inspection points and help with freshness of that unified enforcement. Some vendors in the great reset DLP enforcement layer are also leveraging the Microsoft Security graph for risk signals to enhance threat prevention context coupled with visibility and control of user activities with sensitive data type policies for example conditional access signals like geolocation, endpoint, fingerprints, historical risk scoring,etc.

## Why the browser has become a new enforcement plane

The shift toward web-based user workspaces and the rise of AI chat interfaces have transformed the browser into a vital enforcement plane for data security. As organizations move almost entirely to SaaS applications, new browser-based solutions and extensions are emerging as essential complementary tools to address visibility gaps. While GenAI chat is driving the initial demand for integrated Data Loss Prevention (DLP), the market is rapidly expanding toward agentic platforms that require sophisticated API-integrated services for comprehensive data control. Standard network and endpoint security often fail to capture granular, in-session activities. Consequently, the browser, combined with integrated APIs for agentic services, provides critical last-mile governance over how users and AI agents manipulate sensitive information. Modern browser-oriented security now leverages capabilities like DOM inspection, fingerprinting, and WebAssembly containers to function as a primary enforcement mechanism, offering robust detection, classification, and control at the edge.

Benefits include:

- **SaaS work that happens in-browser:** In today's modern SaaS applications, the most common collaboration and administrative actions occur in web apps.
- **GenAI interactions that happen in-browser:** Chat oriented prompts, copilots, and embedded assistants often run in a browser session even when backed by enterprise models.
- **Copy/paste/upload/download and application context dominate leakage methods in browsers:** This means that last-mile user actions are frequently the decisive exfil path, and they can bypass network or API-only controls without browser and session level enforcement. This capability is achieved either through browser extensions, injections of javascript to monitor dom-tree and execution elements, or enterprise browser replacements.

## The DLP Technology and Deployment Evolution

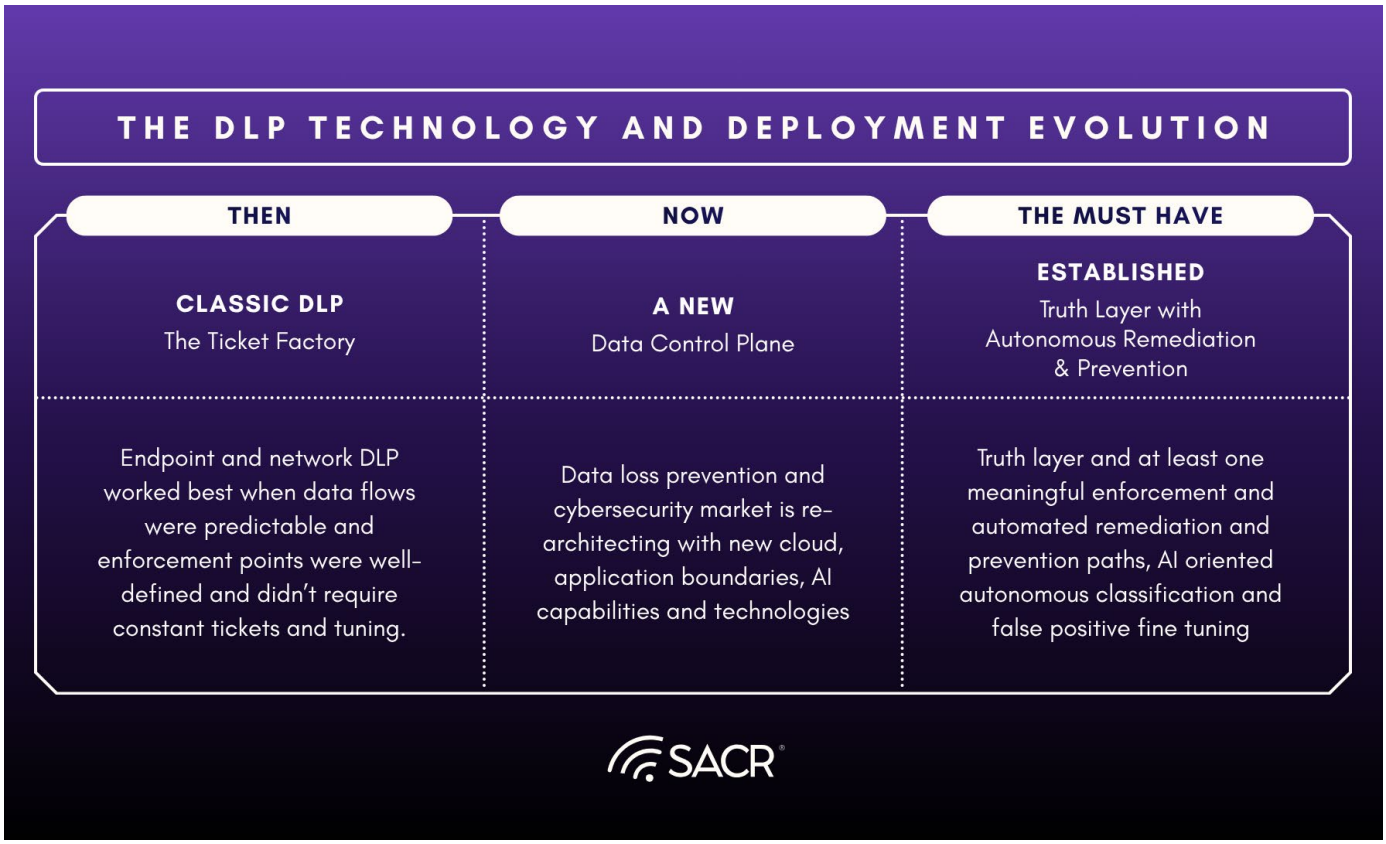

## Modern DLP Outcome Imperatives


**Establish a Discovery-led truth layer** by shifting from Regex to AI Classification, Context Enrichment, and Semantic Understanding to classify data based on actual sensitivity and user intent, eliminating the need for thousands of static regex rules.


**Implement Automated Triage and Streamlined Incident Response** by utilizing Agentic AI to autonomously triage alerts, auto-close false positives, and only escalate validated threats, preventing the SOC from becoming a data loss ticket factory.


**Empower the End-User with Real-Time Just-in-Time Coaching** by leveraging automated workflows and browser-based alerts via platforms like Slack or Microsoft Teams to prompt users for justification or self-correction during policy violations, effectively decentralizing enforcement and reducing low-risk tickets.


**Drive Proactive Remediation and Prevention through Agentic Campaigns** where modern DSPM and DLP solutions evolve toward autonomous, real-time remediation across the truth layer and prevention at the enforcement layer, enabling enterprises to launch agentic communication campaigns or continuous evaluation for prevention and engagement with data owners to address potential risks preemptively.


**Deploy Distributed Automated Enforcement** with machine-speed interventions, including real-time blocking, contextualized DSPM, runtime encryption, and asset quarantine, by deploying uniform protections across endpoints and cloud gateways to effectively halt exfiltration and secure intellectual property.

# When Selecting Data Loss Prevention Controls, Consider the Data Loss Prevention Trade-Off

This infographic illustrates The Data Loss Prevention Value Trade-Off using a four-quadrant matrix that evaluates cybersecurity strategies based on their Value versus their Operational Burden. The chart highlights an Optimal Zone in the upper-left quadrant, where modern solutions like DSPM-DLP convergence, AI-era controls for prompts, and automated remediation and prevention workflows provide high security value with relatively low maintenance effort.

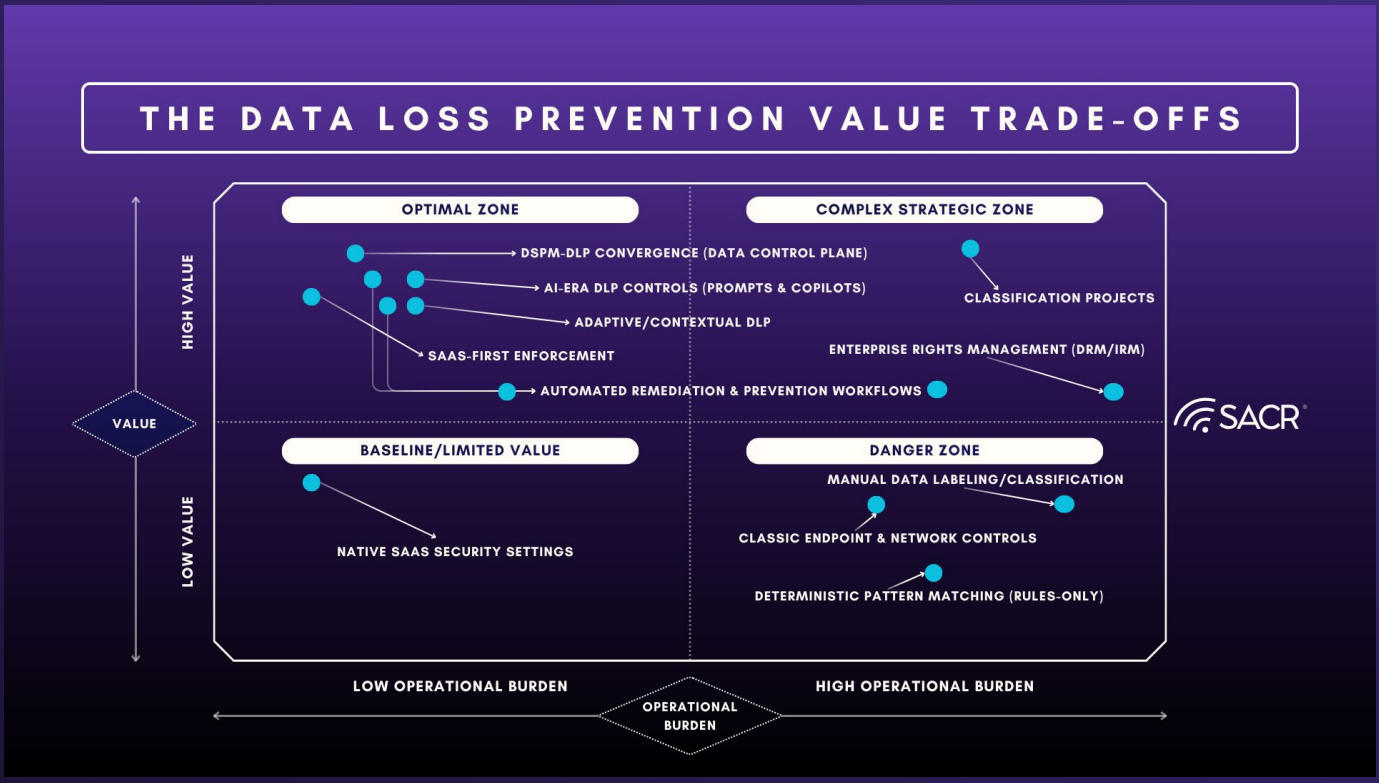


In stark contrast, the Danger Zone in the lower-right quadrant contains legacy methods, such as classic network controls and deterministic pattern matching, which are depicted as having low value and high operational complexity. Ultimately, the visual serves as a strategic roadmap, encouraging organizations to shift away from labor-intensive, rule-based systems toward adaptive, SaaS-first enforcement and automated workflows to maximize efficiency and protection.

## Market Evolution in DLP

# New Building Block Layers Emerge

Modern DLP stacks are increasingly assembled from interoperable component layers rather than a single monolith. This modular approach allows enterprises to move beyond legacy perimeter-based security toward a discovery-led data control plane. The landscape of Data Loss Prevention (DLP) is evolving from static, rule-based systems to highly intelligent, context-aware and AI/agentic enabled platforms designed to secure modern, AI-driven workspaces. Emerging features in this space focus on understanding intent, tracking data throughout its lifecycle, and governing both human and machine identities.

**Here is a look at the emerging concepts and capabilities in modern DLP:**

## Autonomous AI Investigation

Instead of relying on rigid keyword matching or regular expressions that generate massive alert fatigue, modern DLP platforms deploy AI to autonomously investigate potential data loss events. These intelligent systems analyze incidents across multiple dimensions—such as the data itself, the systems involved, human behavior, and the surrounding business processes. By interpreting the context and intent behind a data transaction, the system can distinguish between legitimate business workflows and genuine risks, effectively automating alert triage and response.

## Governance of AI Agents and MCPs

As enterprises increasingly deploy autonomous AI agents to execute tasks, DLP must expand to secure these non-human actors. Emerging platforms provide continuous discovery and governance by enforcing action policies on autonomous behavior. This includes inline inspection of interactions, such as monitoring Model Context Protocol (MCP) calls, to prevent prompt injections, jailbreaks, and unintended data leakage across the AI supply chain. See Agentic Platform examples below.

## Graph-Based Data Lineage

Rather than inspecting files in isolation, modern DLP tracks the entire lifecycle and provenance of data. By capturing a continuous record of where data originated (e.g., a secured internal database), how it has been modified, and who has interacted with it, systems can accurately assess risk. For instance, if sensitive data is copied, reformatted, and pasted into an unauthorized generative AI prompt, the DLP system recognizes the data's sensitive origin and enforces protection policies, drastically reducing false positives.

## In-Line Browser Guardrails

With the web browser becoming the primary interface for SaaS and Generative AI applications, emerging DLP solutions apply granular, last-mile controls directly within web sessions. Instead of simply blocking entire websites, these capabilities monitor text inputs, drag-and-drop actions, and file uploads in real-time. They can dynamically disable copy/paste functions for specific fields or automatically redact sensitive information before it is submitted to public AI models, allowing organizations to adopt AI productivity tools safely.

# Dynamic, Risk-Adaptive Controls

Modern DLP is moving away from static allow or block rules toward dynamic enforcement based on continuous behavioral and intent analysis. By calculating real-time risk scores based on user activity, the system can baseline normal behavior and automatically adjust its posture. If an employee exhibits anomalous or high-risk behavior, such as a sudden spike in downloads or accessing unusual repositories, the DLP controls automatically tighten to intervene. Once the behavior returns to normal, the restrictions are relaxed, minimizing friction for legitimate work.

## Emerging Agentic Platforms in 2026

| Agentic Platform Category | Platform Name |
|---|---|
| **Desktop & Personal OS** | Claude Cowork |
| | Microsoft Copilot |
| | Eigent (Open Source) |
| | OpenClaw |
| **Engineering & Dev** | Claude Code |
| | Verdent |
| | Windsurf |
| | Cline (Open Source) |
| **Business & No-Code** | Vellum AI |
| | Lerty AI |
| | Zapier Central |
| | Airtable Omni |
| **Developer Frameworks** | LangGraph |
| | CrewAI |
| | Microsoft AutoGen |

*Consideration List - not exhaustive

# Converged Concept 1: The Convergence of DSPM and Discovery-Led Truth Layers

- **System of Record:** Discovery-led classification (and label translation and synchronization between Google and Microsoft environments) establishes the definitive truth layer, where Data Security Posture Management (DSPM) acts as the system of record to position DLP as a measurable outcome.
- **New Architectural Model:** DSPM (inventory + classification + risk) → data control plane (orchestration) → distributed enforcement points.

# Converged Concept 2: Adaptive and Contextual Data Control Planes

- **Operational Goal:** Reducing the operational tax and eliminating ticket factory outcomes through high-fidelity signal prioritization and automated triage.
- **Contextual Governance:** Policies must integrate identity, access entitlements, and sharing posture to inform automated remediation and prevention across the full data lifecycle.

# Converged Concept 3: Securing the AI Runtime: Prompts, Copilots, and Agents

- **Leakage Modes:** Addressing emerging exposure paths, including prompt-based exfiltration, Copilot access scope, and agentic tool-call actions across cloud and SaaS surfaces.
- **Runtime Controls:** Implementing session-based interventions—such as redaction, masking, and policy-driven warnings—tailored to specific data sensitivity and user behavior.
- **Auditability:** Establishing governance through audit-grade evidence, capturing data lineage and event provenance to support forensic investigations and remediation or prevention proof.

## The AI-era Example Scenario

**The New Scenario:** An employee pastes proprietary code into a public GenAI chatbot

**Modern DLCP response:** Move towards warning users or redacting sensitive fragments and log an event with audit-grade detail

1. **Technical description:** Prompt inspection, redaction, policy-driven warnings and blocking, Copilot governance, agent tool-call control and logging.
2. **How it addresses the problem:** Covers exposure modes that do not resemble classic file and email exfiltrations.
3. **Integration considerations:** Needs enhanced integrations into GenAI surfaces including but not limited to AI chat interfaces, Developer API Surfaces, Integrated AI services and AI Agent platforms and providers with strong logging and evidence as a differentiator.

# A Deep Dive of the Stack: Modern DLP Layers 0 through 6

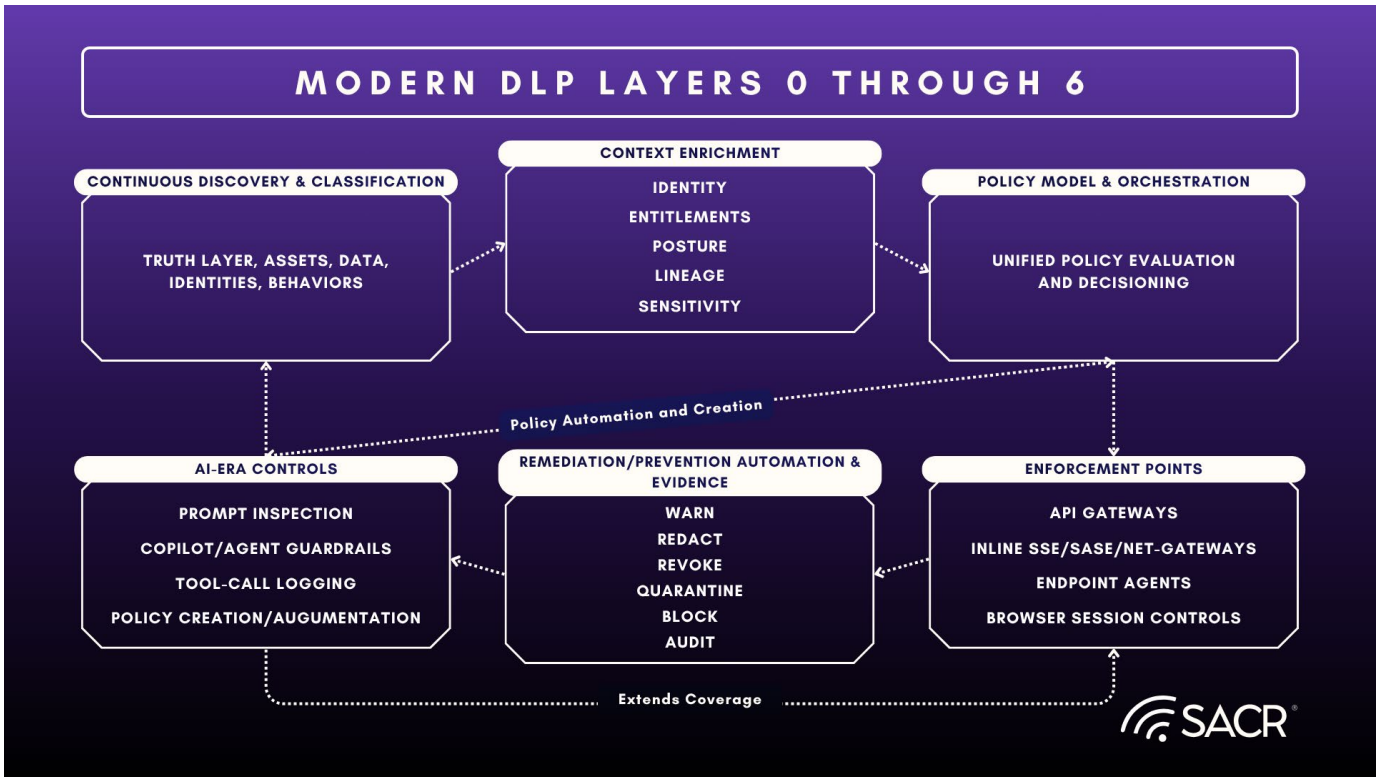


## Intelligence Plane

### Layer 1: Continuous discovery & classification (the truth layer)

- **Technical description:** API-based scanning of SaaS repositories and cloud data stores and performs classification using advanced patterns and machine learning (ML) algorithms with continuous posture updates to reflect real-time changes.
- **How it addresses the problem:** This layer reduces the unknown-unknowns that typically drive the noisy enforcement environments of the past, sporting high false-positive rates, and drawing political pushback from business units. By establishing a definitive truth layer through Data Security Posture Management (DSPM), organizations can position DLP as a measurable outcome of discovery and business enabler or at least an optimizing function.
- **Integration considerations:** Coverage breadth is heavily dependent on the quality of vendor connectors across SaaS, IaaS, and on-premises environments. High-fidelity classification and robust evidence trails are critical for supporting forensic investigations and establishing stakeholder trust in automated actions.
- **Scanning Performance, Cloud Costs and Speed to Discovery:** Various styles of deployment and sampling rates dictate the speed of discovery and scanning. For example, with file scanning using centralized methods, this can impact performance. Scan speed can be enhanced based on localized cloud based scanners and sampling of data stores (especially for structured data types) can be essential for speed (for example, identifying a credit card bin or social security number by sampling a few rows in a database table vs examining all data). Any DLP scanning interactions from Cloud to Cloud or Cloud to Premise can expand cloud costs.

## Layer 2: Access and usage context (identity, entitlement, posture and behavior)

- **Technical description:** Modern DLP can augment content findings with who-accessed and who-shared, privilege and entitlement context, sharing posture, and sometimes lineage.
- **How it addresses the problem:** Applying these contextual elements, especially through federated context facilities such as model context protocol (MCP) reduces false positives and can enhance speed due to reduced re-classification need and enables prioritization or data security controls and enforcement mechanisms (what matters, to whom, and why).
- **Integration considerations:** This layer requires identity and SSO signals, SaaS audit logs, and in some cases endpoint or network gateway/SASE/SSE telemetry.

## Layer 3: Policy orchestration across tools (the data control plane)

- **Technical description:** Ideal deployments will offer the ability to centralize data security policy definitions that map to distributed enforcement points and enforcement templates and apply these suggested policies to reduce manual burden.
- **How it addresses the problem:** Can help prevent policy sprawl across tools and aligns controls to business context and data labeling.
- **Integration considerations:** Orchestration depth varies, many vendors still are reliant on single vendor selection, but some products push policy into SSE and SASE or endpoint, others focus on SaaS actions.

# Enforcement Plane

## Layer 4: Enforcement planes (API, inline, browser/session, endpoint)

- **Technical description:**
  - **API/SaaS-first:** Out-of-band detection + actions (quarantine, redact, revoke links)
  - **Inline SSE/SASE/GW:** Offers real-time inspection/control for web and SaaS traffic (often requires steering and SSL/TLS decryption)
  - **Browser/session (last-mile runtime):** In-session control over copy/paste, upload/download, printing, screen capture, and GenAI interactions
  - **Endpoint:** device-level controls for local exfil paths (removable media, local copies, unmanaged sync clients, print, clipboard)
- **How it addresses the problem:** Provides practical control coverage across the actual runtime surfaces where data moves.
- **Integration considerations:** Inline requires steering, endpoint requires rollout and tuning, browsers may require extension or browser replacement.

## Layer 5: Automated remediation and prevention workflows

- **Technical description:** Automated responses (warn, redact, delete, revoke sharing, label) plus audit trails and investigation context.
- **How it addresses the problem:** Shifts DLP from alert factory to measurable risk reduction and autonomous prevention.
- **Integration considerations:** Action depth is vendor- and connector-dependent, evidence quality is crucial for stakeholder trust.

## Layer 6: AI runtime DLP (prompts, copilots, agents)

### Capability taxonomy: AI runtime DLP

Emerging capability often delivered via API surface or inline proxy runtime) see also SACR Unified Agentic Defense Platforms publication (UADP).

- **Prompt input controls:** Offers paste/upload inspection, warnings, redaction/masking, and policy-driven blocking for sensitive inputs.
- **Output controls:** Offer redaction/masking/ watermarking of generated outputs and safe copy or data export controls.
- **Tool-call governance :** Can constrain what tools can access/send, enforce least-privilege data access, and log tool inputs and outputs where appropriate, controls data residency and trust.
- **Residency and compliance guardrails:** Residency of data and compliance guardrails must exist for the data they inspect and log. This is critical for regulated environments where PII/PHI/PCI data cannot cross geopolitical or operational boundaries, even when processed by an agent or a cloud-hosted tool.
- **Evidence and forensics for agent actions:** Chain-of-custody events for agent activity (who/ what/when), provenance, and investigation-ready timelines.
- **Technical description:** Prompt inspection, redaction, policy-driven warnings/blocks, Copilot governance, agent tool-call logging.
- **How it addresses the problem:** Covers exposure modes that do not resemble classic file/email exfiltration.
- **Integration considerations:** Needs integration into GenAI surfaces, strong logging and evidence is a differentiator.

## Enterprise Data Loss Prevention Great DLP Reset Deployment Framework (Ideal Scenario)

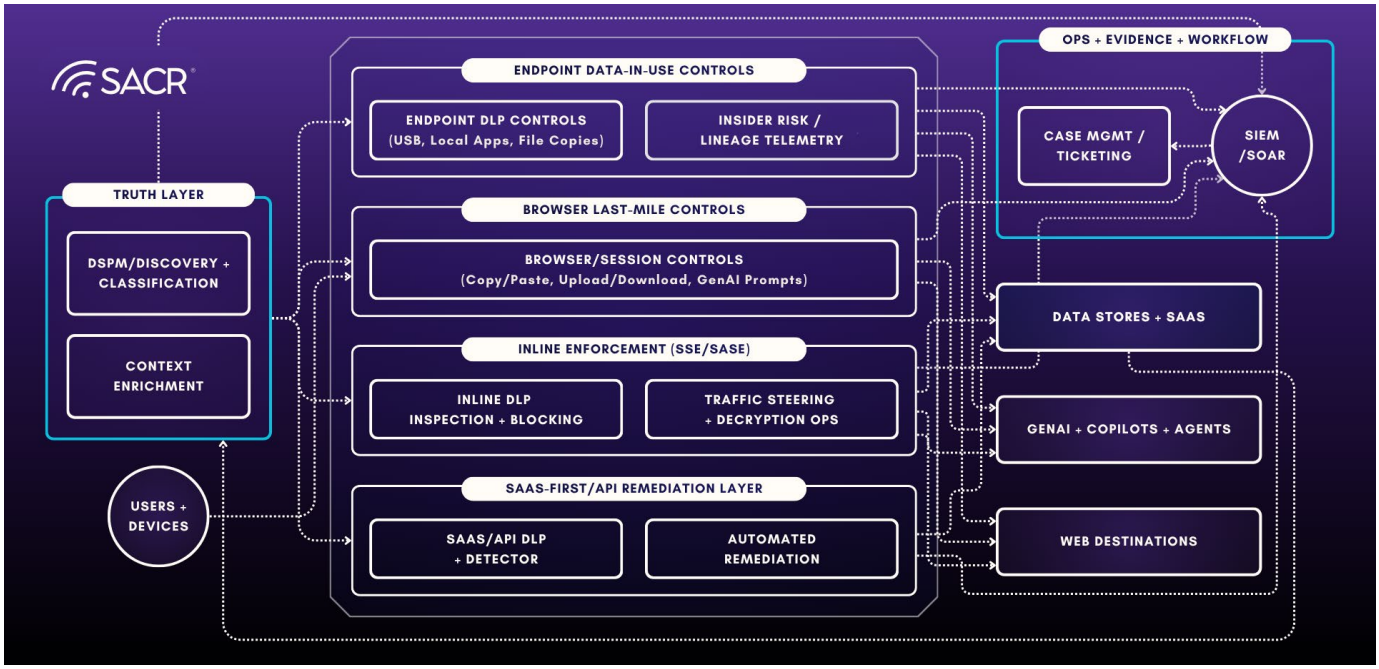

# Market landscape: DLP layers (archetypes)

This section maps common product layer archetypes to the control-plane model. Many platforms span multiple layers, but most have a primary center of gravity.

| DLP Layers | Description | Center of Gravity |
|---|---|---|
| **Truth layer / DSPM-led intelligence plane** | Continuous discovery and classification, posture, and risk context, which feeds enforcement and remediation. | Feeds enforcement and remediation and prevention actions. |
| **API/SaaS-first** | Agentless/API-led policy and actions inside collaboration applications; strongest for fast outcomes in SaaS oversharing. | Fastest outcomes in SaaS oversharing. |
| **Inline SSE/SASE/GW** | Network-path enforcement for real-time web and SaaS controls; provides power but carries a higher architecture and tuning burden. | Real-time web and SaaS controls. |
| **Endpoint** | Endpoint DLP and behavioral/insider controls for local exfil paths. | Local exfiltration paths. |
| **Browser/Session (Emerging)** | Enterprise browsers and extensions delivering in-session DLP for SaaS and GenAI workflows. | In-session DLP for SaaS and GenAI workflows. |
| **AI runtime DLP (Emerging)** | Prompt/output/tool-call governance, agent guardrails, and audit-grade agent evidence. | AI prompt, output, and agent governance. |

# Adoption path (maturity model)

1. **Start with the truth layer:** Can establish discovery and classification and basic evidence.

2. **Prove remediation and prevention in one high-noise channel:** Pick a SaaS surface and close the loop with reversible actions.

3. **Add context to reduce noise:** Add context from sources such as identity,entitlements and sharing posture to prioritize what matters most for preventive enforcement actions or remediation.

4. **Add heavier enforcement intentionally:** Inline SSE/SASE/GW (and SSL/TLS decryption) and endpoint where required by risk scenarios.

5. **Make AI runtime governance explicit:** Emerging prompt,output, tool-call controls offer enforcement of DLP policy plus may add AI agent evidence through MCP monitoring or API integrations. Treat browser and sessions as a primary runtime where applicable. Also see SACR publication (Unified Agentic Defense Platforms (UADP))

6. **To help Unify Enforcement Policies:** Develop and use Normalized Data Sensitivity Labeling across Enforcement Points. Data sensitivity labeling from DSPM serves as a ground truth layer, derived from Microsoft Purview and Google Workspace classification label schemes.

## Data Sensitivity Labeling:

# Normalized for Enforcement Layer Policies

| Level | Normalized Label | Risk Profile | Microsoft Purview Equivalent | Google Workspace Equivalent |
|---|---|---|---|---|
| 0 | **Public** | No Risk; intended for public consumption | **Public** | **Public** |
| 1 | **General** | Standard business data; no PII | **General** | **Internal** |
| 2 | **Confidential** | Sensitive business info or PII/GDPR data | **Confidential** | **Confidential** |
| 3 | **Restricted** | Intellectual Property or Sensitive PII | **Highly Confidential** | **Restricted** |
| 4 | **Top Secret** | Critical business survival (M&A, Keys). | **Restricted - [Project]** | **Secret / Private** |

SACR®

# What DLP Vendors Should do to Win in the Great DLP Reset

The DLP market is being re-oriented by buyers against a new baseline: measurable risk reduction with lower operational burden, across modern runtimes (SaaS, cloud data platforms, and GenAI). Vendors that continue to lead with more detections without proving enforceable outcomes will increasingly be treated as noise generators rather than control-plane platforms.

## 1) Make operational burden a first-class product outcome

Buyers now treat deployment complexity, policy sprawl, and false-positive triage as existential program risks. Vendors should:

- Ship opinionated defaults (starter policies, templates, and tuning guardrails) rather than assuming every customer will build a program from scratch.
- Provide staged rollout and rollback mechanics that are predictable and safe (preflight checks, safe modes, clear blast-radius controls).
- Instrument and report burden: time-to-deploy, time-to-first-signal, false-positive rate, triage hours/week, and disruption rate (how often enforcement breaks legitimate work).

## 2) Prove time-to-first-signal and prevention in days or hours, not quarters

A major separation in the market is how fast a platform can surface "material risk" (not just matches). Vendors should design onboarding around a 1–2 week proof window, sooner if at all possible:

- Fast connectors to core data gravity (M365/ Google, Slack, Salesforce, GitHub, key cloud stores).
- Immediate prioritization (what's sensitive, who can access it, what's externally exposed).
- At least one closed-loop remediation and prevention enforcement paths early (revoke public links, quarantine, ticket, redact, block, delete) so visibility becomes risk reduction.

## 3) Turn context and LMM natural language into a real differentiation lever

Modern DLP decisions are increasingly identity and entitlement-driven. Vendors should operationalize context:

- Identity, entitlements and sharing posture should directly reduce noise and improve prioritization.
- Explanations must be human-usable: why this object mattered, why this actor and action is risky, and what changed after remediation.
- Enhancing Insider Threat incidents with natural language cognitive Large Language Model (LLM) outputs can significantly increase storyline elaboration on incidents and events.

## 4) Be explicit about enforcement points, and avoid one-control-point narratives

Buyers are increasingly skeptical of vendors that imply universal coverage from a single enforcement surface. Vendors should clearly articulate:

- Where controls actually execute (SaaS/API, inline SSE/SASE/GW, endpoint, browser/session, email).
- Which actions are enforceable per channel (block, quarantine, revoke sharing, redact/ delete, label/classify, coach/warn, ticket/ workflow).

- How policy intent stays consistent across distributed control points (a real control plane vs. disconnected features).
- How they properly stitch together events and data from various sources across the DLP overall deployment architecture to properly create incidents and perform remediation actions in-context, or nudges to users without creating fatigue.

## 5) Win on remediation depth (with guardrails), not alert volume

The market is shifting from find to fix, nudge,inform. Vendors should strengthen remediation and prevention depth and safety and their ability to lightly engage users:

- Prioritize reversible and low-friction actions first (revoke sharing, quarantine, remove public access) before heavy blocking.
- Provide automation with guardrails (approvals, exception handling, rollback, and proof-of-remediation and prevention).
- Track outcomes that matter: exposure reduction, % auto-remediated, MTTR, and repeat-offender reduction.
- Engage users lightly through agentic communications via communications channels (Slack, Teams, etc) for light nudge and user contextual education.

## 6) Treat GenAI and agent platforms as a default runtime

GenAI and agent platforms introduce high-frequency leakage paths (prompts, uploads, outputs) and emerging MCP agent/tool-call surfaces. Vendors should:

- Ship concrete prompt/output controls (detect, warn, redact, block) based on sensitivity and context.
- Provide audit-grade logging for GenAI interactions, including relevant inputs/outputs where feasible and permitted.
- Package GenAI policies as templates aligned to real data types (source code, credentials/ secrets, regulated identifiers, contracts/M&A).
- Consider and Expand integrations and capabilities towards Agentic Workflow and agentic platforms (See consideration list for supported data loss controls in the mapping below)

## 7) Raise the evidence bar with: audit-grade, investigation-ready artifacts

As DLP becomes a control plane, evidence and chain-of-custody become competitive wedges. Vendors should:

- Attach defensible evidence to each high-impact event: actor, object, action, sensitivity, destination, timestamps, and remediation or prevention result.
- Provide investigation-ready timelines and, where possible, lineage/provenance signals that support incident response and audit readiness.

## 8) Leverage standardized labeling, and seek cross-product integrated enforcement

Since customers increasingly utilize various vendors in their data loss prevention and data security programs, it's incumbent on existing vendors to work more harmoniously together, sharing enough information properly to improve the enforcement and control layers. Vendors should:

- Integrate or develop sharing mechanisms between discovery and control planes to properly utilize standardized labeling schemes if this capability is not already present.
- Consider leveraging the Microsoft security graph (as an example) and other sources of risk information as common context sources for enforcement or as elevated risk signals

# Market Competitors:
# Data Loss Prevention (DLP)

**Key Vendor Differentiators:**

Automation depth, Audit-grade evidence, Data Lineage, Closed-loop Autonomous Tuning

**Great Reset Vendor Alignment Archetypes delivering Modern DLP capabilities:**

- API / SaaS-first
- Inline SSE/ SASE/GW
- Endpoint agent
- Browser extensions
- Enterprise browsers
- DSPM-led control plane
- Hybrid


ORION

# Orion Security

## Vendor Profile

Orion Security is an AI-native data loss prevention platform positioned to reduce the operational brittleness of traditional DLP by shifting emphasis from static, manually maintained policies to context-rich detection of why sensitive data is moving. Its center of gravity is real-time detection and prevention of risky data movement across multiple enforcement points (endpoint agent, browser extension,email gateway, and SaaS/API integrations), with particular focus on high-noise environments where security teams struggle to distinguish legitimate workflow activity from true exfiltration, especially as GenAI usage increases the volume of sensitive data interactions.

## Products/Services Overview

- **Unified DLP platform spanning multiple enforcement surfaces**
  - Endpoint agent capabilities intended to observe and control local data actions and common exfiltration paths.
  - Browser extension capabilities intended to control last-mile user actions in web/SaaS and GenAI interactions.
  - Email Gateway that adds another security layer to prevent data exfiltration though emails, even when sent through unmanaged devices (like mobile phones).
  - SaaS/API integrations intended to observe and act on risky data movement and oversharing in connected services (action depth varies by connector and must be validated).
- **Context- and behavior-oriented detection engine (“beyond policies” positioning)**
  - Uses context signals (identity, destination, environment, and lineage) to judge likelihood of data loss vs normal business activity.
- **Data lineage / tracing-oriented investigation support**
  - Emphasis on mapping or reconstructing data movement paths to improve triage and provide investigation narrative (what happened and how data moved).
- **User interaction and enforcement workflows**
  - Vendor materials describe controls such as warnings/coaching, justification/override patterns, and automated blocking, with escalation options.
- **Integration into security operations workflows**
  - Vendor-provided materials describe integrating into existing SOC/case management patterns.

## Market Category

Enterprise DLP Suite

## Great DLP Reset alignment

Hybrid

Orion aligns with the Great DLP Reset thesis by treating DLP as a context-driven control plane paired with distributed enforcement, rather than a policy spreadsheet attached to a few chokepoints. The model is oriented to using richer context to reduce false positives and avoid policy sprawl, placing controls closer to where modern data movement occurs (endpoint and browser-mediated SaaS and GenAI workflows), and providing a more investigation-ready view of data movement (lineage/flow framing). The primary tradeoffs are typical of emerging AI-native DLP approaches: buyers must validate explainability and governance of AI-driven decisions, confirm the practical breadth of integrations and enforcement depth, and ensure deployment/change-management (agents and extensions) is acceptable at scale.

## Overall Viability and Execution

Public reporting indicates Orion raised a significant funding round in early 2026, and investor commentary portrays strong early go-to-market momentum. This supports near-term viability and suggests the company is investing aggressively in product development and enterprise sales, but it does not substitute for customer diligence on support maturity, roadmap stability, and implementation outcomes across diverse environments.

Orion is described as highly responsive during RFP/POC phases, with a set-and-forget aspiration (reduced tuning burden versus legacy DLP). What tends to go well are deployments, fast initial visibility and a clearer signal-to-noise story when buyers are already suffering from alert fatigue. What tends to be hard: enterprise-scale rollout (endpoint + browser), providing consistent enforcement across varied data channels, and meeting the expectations of organizations that require mature global support, deep compliance artifacts, and highly deterministic controls for regulated workflows.

## Core Functions and Use Cases

- **Reducing false positives and policy sprawl in DLP programs**
  - Use context/behavior to improve precision and reduce ongoing tuning overhead.
- **Securing SaaS and browser-mediated workflows (including GenAI use)**
  - Control and monitor common last-mile leak paths such as web uploads, copy/paste, and prompt entry.
- **Endpoint-centric prevention for data-in-use actions**
  - Address local exfil paths (e.g., file movement patterns, local application interaction) with agent-based controls.
- **Incident investigation and evidence-building for data movement**
  - Provide lineage/flow visibility to accelerate triage and support auditability.
- **Augmenting existing DLP investments rather than forcing rip-and-replace**
  - Public interview content indicates Orion may run alongside incumbent stacks in larger enterprises, focusing on context enrichment and reduction of noisy detections.

## Use Cases and Pain Points Addressed

1. **Cutting noise in existing DLP deployments without weakening enforcement**
   - Enabling capability: context-aware detection to distinguish legitimate business workflows from suspicious movement.
   - Why it matters: reduces analyst fatigue and increases trust in DLP signals.
2. **Preventing sensitive data entry into public GenAI tools via browser sessions or desktop apps**
   - Enabling capability: Browser extension and enforcement for copy/paste and prompt submission patterns desktop app for data exfiltration prevention.
   - Why it matters: GenAI creates a high-frequency, human-driven leak path that traditional DLP often misses or over-blocks.
3. **Stopping risky uploads/shares from endpoints to SaaS destinations**
   - Enabling capability: endpoint agent + SaaS context to judge destination risk and take actions (warn/block/justify).
   - Why it matters: many leaks are "normal" user workflows pointed at the wrong destination or account context.
4. **Detecting anomalous exfiltration behavior (insider or compromised identity patterns)**
   - Enabling capability: behavior/intent analysis combined with identity and environment signals.
   - Why it matters: material incidents often look like legitimate access until correlated with context and unusual movement.
5. **Building investigation-ready narratives of how data moved**
   - Enabling capability: lineage/flow mapping to show the sequence of movement and implicated users/apps.
   - Why it matters: speeds containment decisions and improves defensibility in audits and post-incident reviews.

## Differentiation and Competitive Novelty

- **Beyond policies positioning: intent and workflow-aware DLP rather than rule-first DLP**
  - **Competitive context:** contrasts with legacy suites that rely heavily on static patterns and long tuning cycles.
- **Multi-surface enforcement portfolio (endpoint agent, browser extension and SaaS/API)**
  - **Competitive context:** attempts to unify prevention across the most common modern leak points without requiring a full SSE/SASE-inline architecture.
- **Lineage/flow emphasis as a primary mechanism for trust and triage**
  - **Competitive context:** aims to compete on evidence quality and analyst usability, not only detection.
- **Coexistence model for large enterprises**
  - **Competitive context:** public interview material indicates Orion may supplement entrenched platforms (e.g., to reduce noise and add context) rather than demanding immediate replacement.

## SACR Key take away:

Orion Security is best suited for CISOs who are dissatisfied with legacy DLP's tuning burden and false-positive fatigue and want a more context-driven approach to preventing real data loss, especially in SaaS-heavy environments where GenAI usage and browser-mediated workflows dominate. Shortlist Orion when you need faster, higher-trust signal and practical enforcement at endpoint and browser leak points, and when the organization is willing to evaluate an emerging vendor's AI-driven detection model through a rigorous pilot.

# Practical Recommendations for CISO's and Practitioners

## Buyer Starting Point Decision Tree

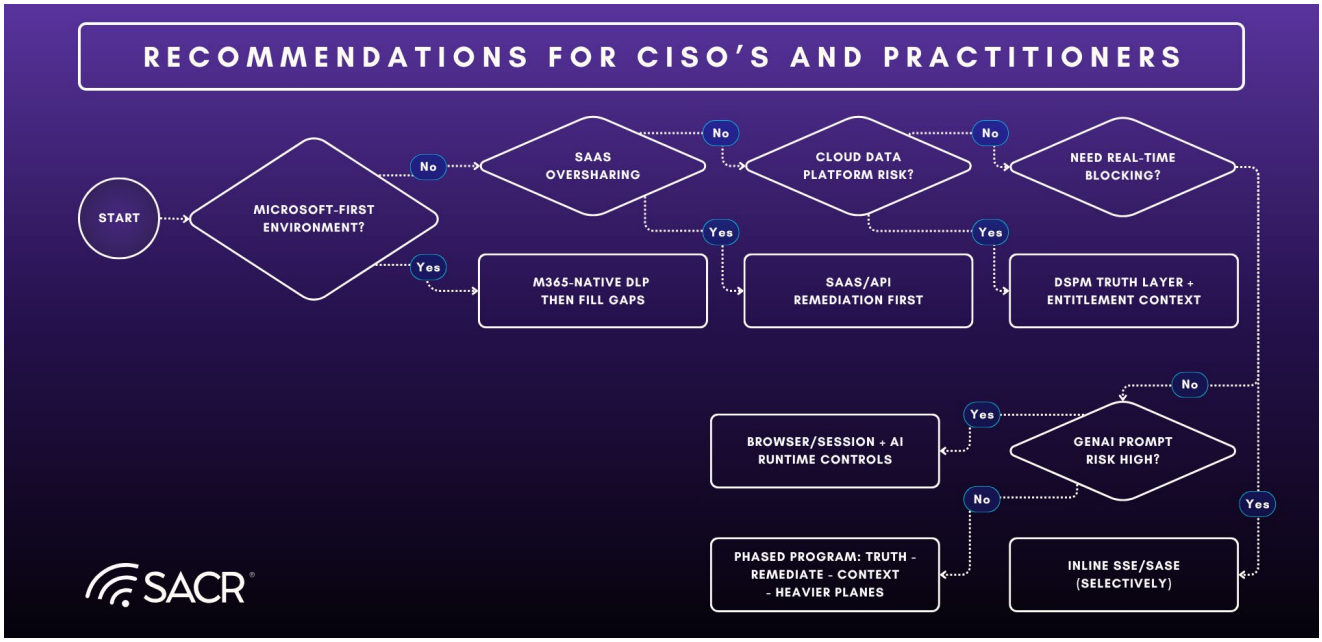


DLP Reset Program Focus Recommendations by phase

- **Phase 0 (classic):** Retain for regulated identifiers and mature endpoint and network needs, but constrain scope and staffing expectations, refresh and integrate to advance to later phases.
- **Phase 1 (SaaS and cloud sprawl):** Prioritize SaaS-first enforcement and remediation in the noisiest collaboration channels.
- **Phase 2 (control plane):** Invest in discovery and classification and the addition of identity and entitlement context, then focus on automated remediation and workflows where possible, or human in loop for exception handling or deeper triage and remediation validation (if needed).
- **Phase 3 (AI-era):** Implement explicit prompt and agent controls and browser and session last-mile controls with strong logging and ideally mapping data lineage from end (source) to end (destination).

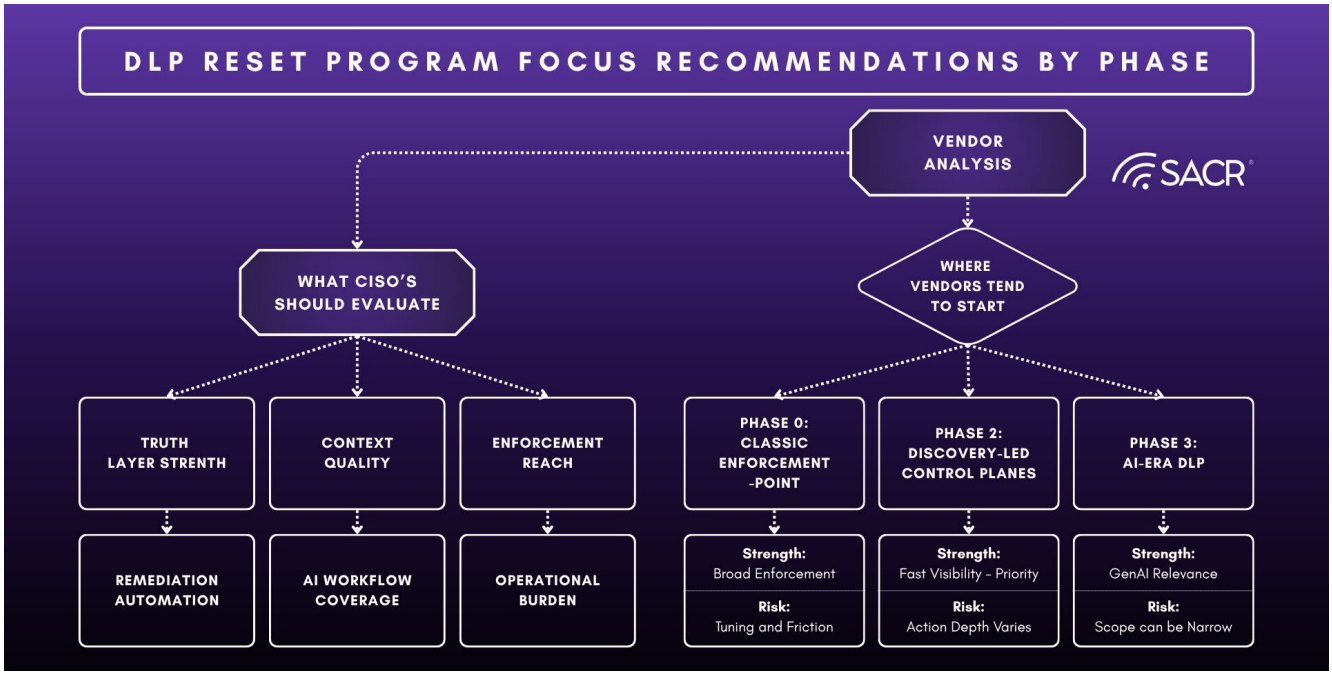

# Actionable Security Program Steps for CISOs and Security Leaders

## 1. Establish the truth layer first (discovery + classification)

- **Implementation considerations:** prioritize highest-risk SaaS and cloud data stores, define sensitivity taxonomy.
- **Success metrics:** % of sensitive data discovered and classified, time-to-first-signal.
- **Timeline:** Weeks for initial coverage; ongoing for expansion.

## 2. Map your enforcement surfaces to real runtime environments

- **Implementation considerations:** Decide where API actions suffice vs where inline SSE or endpoint is necessary.
- **Success metrics:** % of high-risk channels covered by enforceable controls.
- **Timeline:** 1–2 quarters depending on steering and endpoint roll out.

## 3. Reduce policy burden with context-rich decisions

- **Implementation considerations:** Integrate identity, entitlements, sharing posture, and behavioral signals to manage insider risk and reduce false positives in triage and policy.
- **Success metrics:** False positive rate, analyst time per incident.
- **Timeline:** Incremental, measurable within 30–60 days post integration.

## 4. Prioritize automated remediation and prevention over alerting

- **Implementation considerations:** start with reversible actions (warn and revoke link), then escalate to redact,delete,quarantine.
- **Success metrics:**

1. Mean time to remediate (MTTR) and mean time to prevention
2. % incidents auto-remediated
3. Business disruption rate
4. Timeline: 60–120 days to mature workflows.

## 5. Make AI workflow governance explicit

- **Implementation considerations:** define controls for prompt pasting, Copilot scopes, and agent tool-call logging.
- **Success metrics:**

1. % GenAI apps covered
2. Number of policy-enforced AI events
3. Audit completeness
4. Timeline: 30–90 days for initial controls depending on surfaces.

## 6. Treat browser and session controls as last-mile DLP where needed

- **Implementation considerations:** Use targeted deployment for high-risk groups and workflows.
- **Success metrics:** reduction in SaaS screenshot, copy, download leakage events.
- **Timeline:** pilot in weeks; expand by cohort.

## 7. Align the program to measurable outcomes (e.g., compressing MTTR and anchoring the program in quantifiable risk reduction outcomes)

- **Implementation considerations:** Define risk scenarios (exposed data in SaaS, oversharing to copilots, public repo leaks).
- **Success metrics:** Reduction lower mean time to resolution (MTTR) in exposed sensitive objects, fewer critical exposures, improved evidence trails. (e.g., *reduce MTTR*, *align to measurable outcomes*)
- **Timeline:** Ongoing, baseline within a quarter.

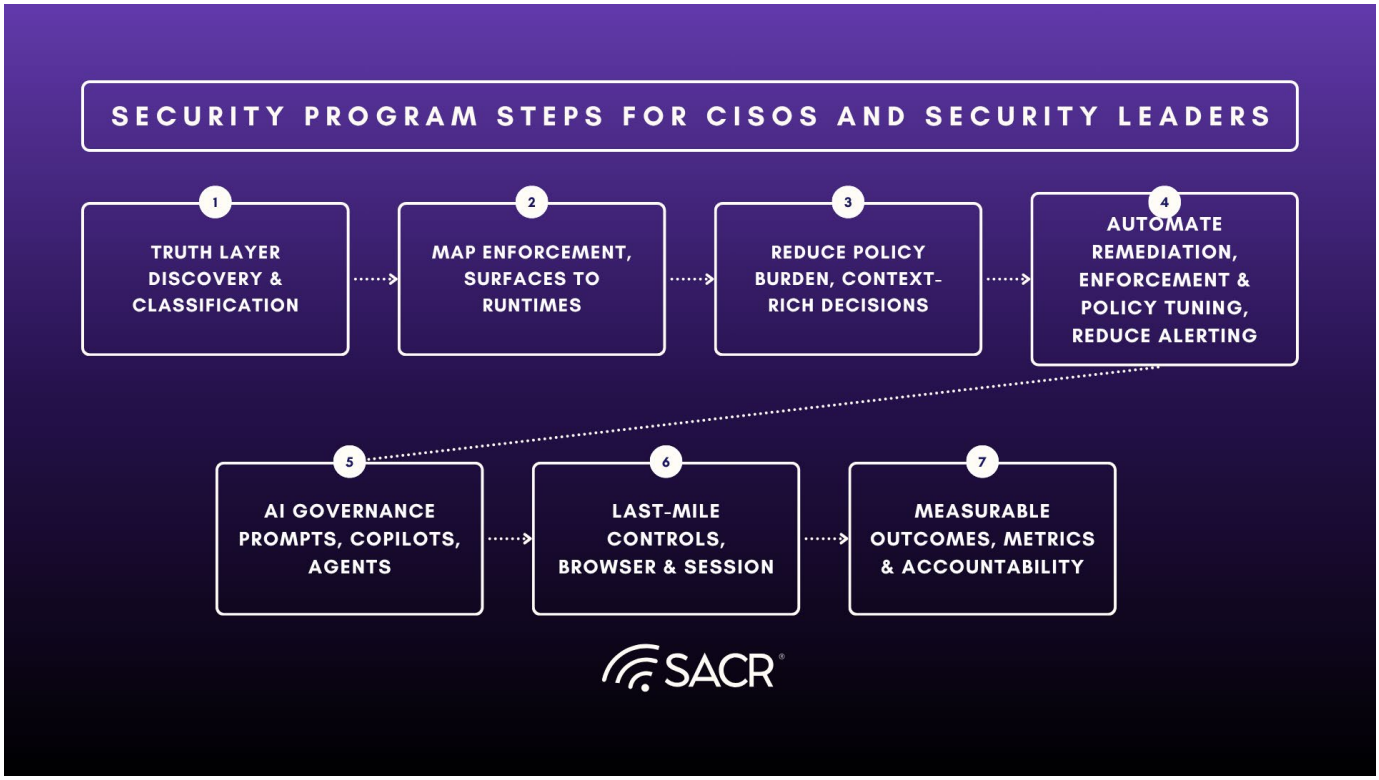


## Security Engineering and Architectural Practitioner Guidance

**First:** Establish the truth layer, then prove remediation in one high-noise channel.

Start by building a defensible inventory of sensitive data (discovery + classification), then pick a single collaboration channel where risk is visible and operational friction is high (e.g., Drive and SharePoint, Slack and Teams, GitHub) and close the loop with a small set of reversible, automated actions.

- **Scope:** 1–3 repositories and apps with the most sensitive data and the most frequent oversharing and exposure patterns.
- **Controls to emphasize:** Classification and evidence trails first; automation second (revoke external links, quarantine, label, owner notification).
- **Success criteria:** Faster time-to-first-signal, measurable reduction in exposed sensitive objects, and a remediation workflow that doesn't become a ticket factory.
- **Avoid:** Starting with broad block policies or multi-channel rollouts before you have trusted classification and stable workflows.

**Next:** Enrich findings with identity and entitlements context and expand repository coverage.

Once you can find sensitive data and take consistent action in one channel, add the context required to prioritize what truly matters (who has access, how it was shared, what level of exposure exists) and broaden coverage to the next set of repositories.

- **Add context:** identity signals (SSO and IdP), group membership, entitlement and access graphs, sharing posture, and (where possible) lineage and audit context.
- **Expand systematically:** Add the next repositories based on risk scenarios (customer data stores, executive collaboration spaces, developer ecosystems).
- **Improve triage:** Drive down false positives and analyst time-per-incident by prioritizing high impact and high exposure findings.
- **Standardize policy intent:** Keep policies consistent as coverage grows (same sensitivity taxonomy, same response ladder).

**Then:** introduce heavier enforcement points only where the risk scenario demands it (inline SSE and endpoint).

Inline and endpoint controls are powerful, but they introduce architectural and operational tax. Add them after discovery-led visibility and remediation are working, and only for scenarios where out-of-band and API actions are insufficient.

- **Use inline SSE when:** real-time control is mandatory (regulated egress), you need web and SaaS traffic inspection, or you must prevent exfil in-session.
- **Use endpoint controls when:** device-level actions are the dominant exposure path (removable media, local copies, un-managed sync clients, print and screen capture).
- **Pilot with a narrow cohort:** high-risk teams and users, a small set of apps, and a clear warn and block escalation path.
- **Keep the program measurable:** track tuning effort, user friction, and risk reduction so enforcement doesn't recreate the classic DLP burden.

## Future View: Emerging DLP Vendors and Trends

This section shifts focus to the emerging vendors and new architectural narratives that are shaping the next generation of Data Loss Prevention (DLP). It moves beyond established players to examine startups and specialized platforms that prioritize AI-driven intelligence, rapid time-to-value, and addressing specific high-friction challenges like insider risk and SaaS collaboration leakage. These companies represent the cutting edge of the DLP Reset, often employing agentless models, edge AI, and deeply integrated remediation workflows to deliver measurable risk reduction with minimal operational burden, signaling the future direction of data security control planes.

# The Great DLP Reset AI and Agentic Transformation

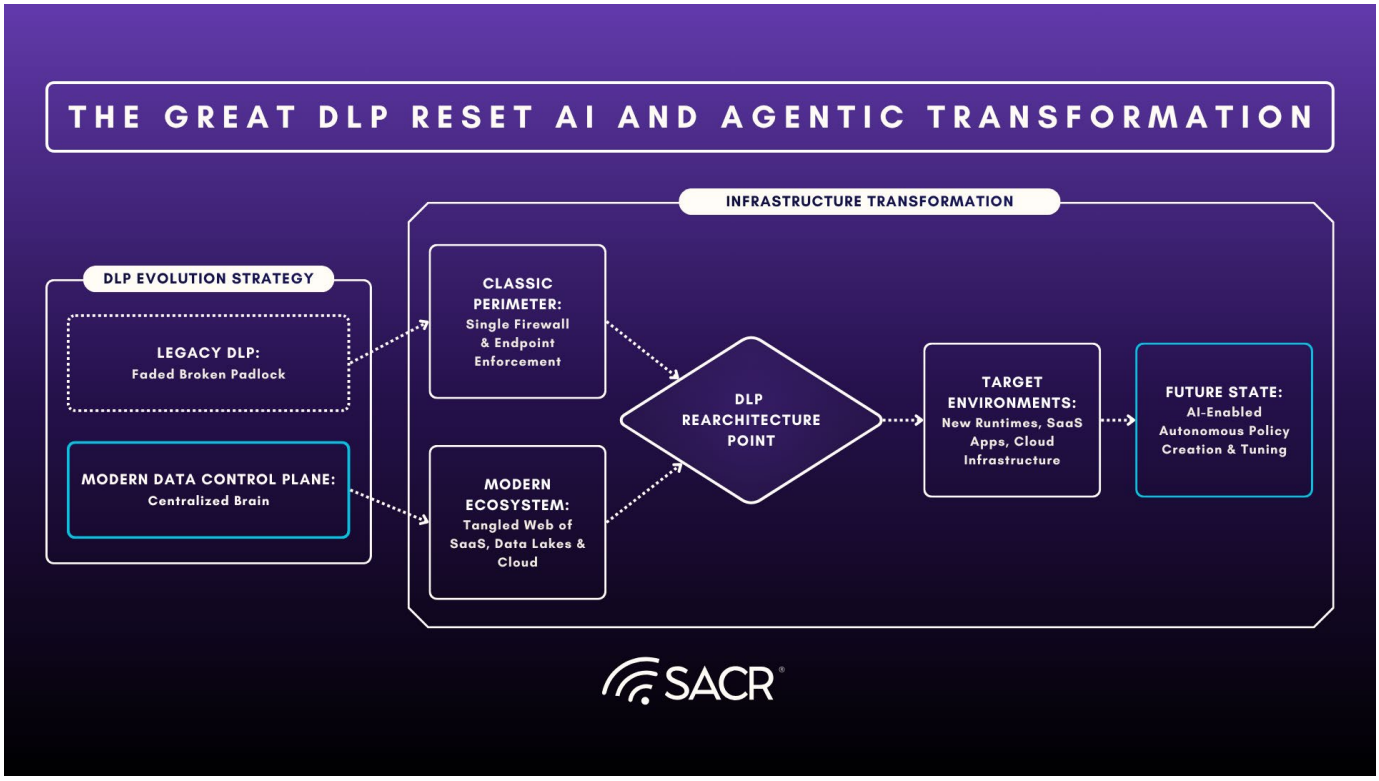

# Normalized Features Across DLP Vendors (Total of 42 vendors Assessed)

The following chart depicts the features across the in-scope Data Loss Prevention providers (total of 42) in SACR exploration of DLP vendors in the market in 2026. The features emphasize compliance oriented mandates such as providing evidence and auditability, classification of data, policy and orchestration features with the fourth priority being automation and remediation.

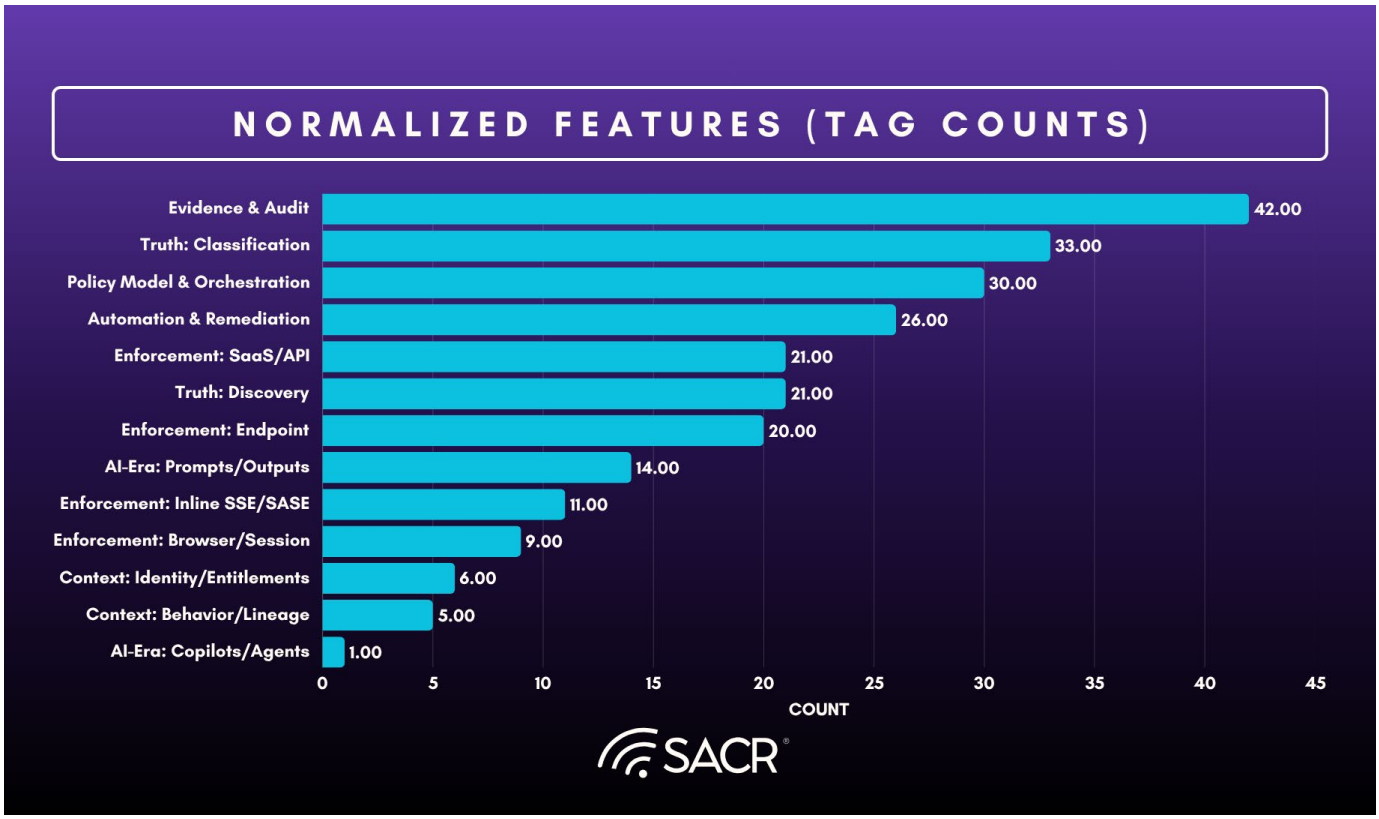


Below is a list of key differentiating features SACR found across data loss prevention vendors. The lower end of the features are either emerging capabilities (e.g. net new features) or features not prioritized by all vendors in the market. Features including operations oriented features such as endpoint rollout capabilities, compliance focused features such as templates and compliance packs and ability to perform revoke of sharing links and the ability to perform redact and delete. Emerging areas of stronger client interest for example in AI-era features such as prompt redaction and LLM semantic extraction from content are aligned to our called out AI-era focus. Also noteworthy is JIT coaching (just in time coaching) for guiding users in real-time to tell them how they ought to handle data and auto-tuning, also aligned to our view of future features in the data loss prevention (DLP) market.

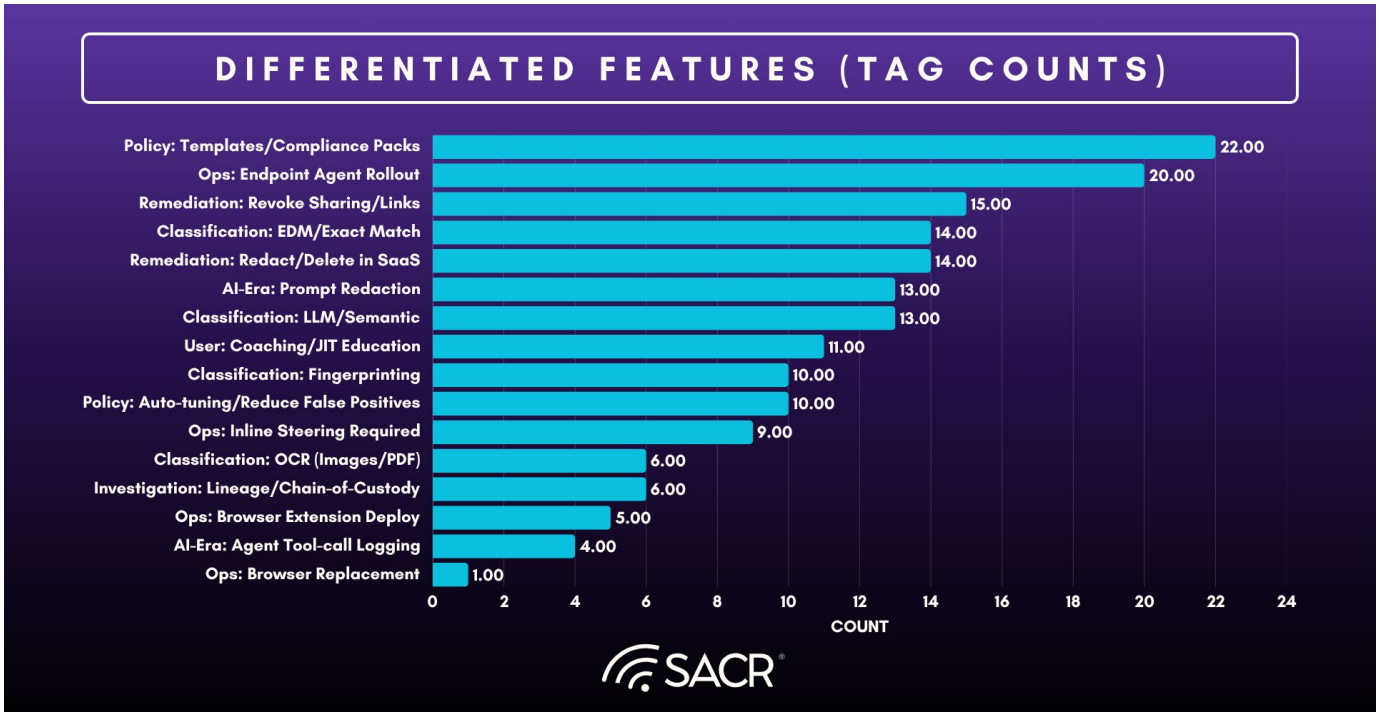


## AI-era DLP is Causing Changes Across Several Key Areas

DLP is critical given the massive expansion of sensitive data across SaaS platforms, cloud, AI, AI agents and premises environments, and into centralized repositories like data lakes, all of which present new and complex vectors for data loss. AI is increasingly being used in DLP solutions and being added to the complex architectures used, while also enabling new capabilities in emerging DLP solution features.

Below is a chart of new emerging classification technologies and their penetration and use by various vendors in our DLP Reset market analysis. Notice that technologies like LLM-based classification and ML semantics have emerged but are lower than traditional rules/regex/pattern style classifiers. This is because these are emerging technologies and capabilities in the latest generation of AI-era tools. We also noted that image recognition, OCR (Optical Character Recognition) and PDF examination was another area of variance between vendors that is often unsupported in some vendors and enforcement channels.

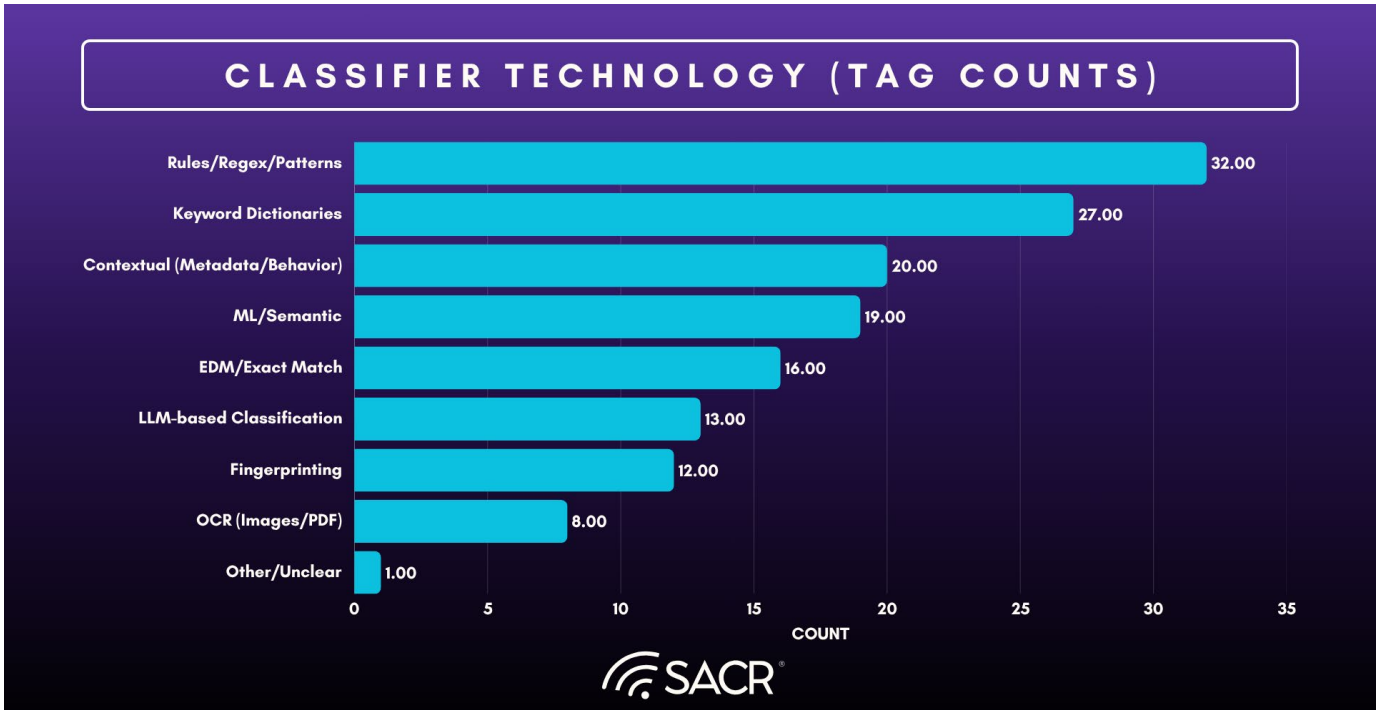

### AI-era DLP is evolving in several key areas:

- **LLM/semantic classification as the new detection engine:** Higher-fidelity classification of unstructured data (docs, chat, code) and fewer brittle regex-only policies.
- **Truth-layer + context becomes mandatory:** Discovery-led inventories (DSPM-style) plus identity, entitlements, and sharing posture help AI and agents decide and take action on what is materially risky.
- **GenAI prompt and output controls (narrow but urgent):** Controls for prompt pasting and uploads, output redaction and masking, and policy-driven warnings and blocks in AI chat and copilots.
- **AI and Agent governance expands the DLP surface area:** Tool-call logging via MCP and platform integrations with Co-Worker and agentic platforms, least-privilege data access for agents, and guardrails that constrain what agents can retrieve/send.
- **Browser and session becomes a first-class enforcement plane:** Last-mile controls over copy/paste, upload and download, printing, and screen capture across SaaS and GenAI workflows.
- **Shift from alerting to automated remediation:** Revoke links, quarantine, redact/delete, label, and other actions that reduce exposure without creating ticket-factory operations.
- **Audit-grade evidence as a competitive wedge:** Investigation-ready timelines (actor, object, action, timestamp, remediation proof) and lineage and provenance where feasible.

### Market bifurcation in the near term:

- **Guardrail point solutions** (prompt, agent and runtime visibility and control) with high relevance but narrower breadth.
- **Broader DLP and control-plane platforms** claiming agentic autonomy and cross-plane orchestration have high upside, but require proof of real enforcement depth and measurable burden reduction.

**Longer term expected convergence towards Unified Agentic Defense Platforms** (see SACR UADP report)

## Emerging DLP Vendors to Watch

This section shifts focus to the emerging vendors and specialized platforms that are shaping the next generation of Data Loss Prevention (DLP). These startups prioritize AI-driven intelligence, rapid time-to-first-signal, and addressing high-friction challenges like insider risk and SaaS collaboration and data leakage. By employing agentless models, edge AI, and deeply integrated remediation workflows, these companies deliver measurable risk reduction with minimal operational burden, signaling a future direction of Unified Data Control Planes (UDCP) and control of data (as well as other cybersecurity functions) in the browser or on endpoints.

# Conclusions

DLP is being rebuilt into a new Unified Data Loss Control Plane (DLCP) because enterprise data no longer lives behind a few enforceable chokepoints, data at rest is difficult to map and share with other data security tooling, AI is leaking data through shadow AI and AI workflows have introduced entirely new leakage paths through coming agent to agent interactions. The market direction is clear: discovery-led truth layers, context-rich policy decisions, and automation-driven remediation are replacing the old model of static rules and endless tuning. Solutions are moving to runtime enforcement vs detection and response through manual processes, and autonomous operations are becoming available rapidly to enhance data defense in the era of realtime agents and greater sprawl.

**SACR Key Takeaway:** For CISOs, the strategic implication is to treat DLP as a control plane program rather than a single product. Winning strategies start by making sensitive data and access realities visible, then apply the minimum necessary enforcement across SaaS APIs, inline controls, endpoints, and browser and AI surfaces, with remediation automation and evidence quality as the core success measures. DLP is dead. Long live the Data Control Plane. The DLP reset favors platforms that enable a data control plane and programs that reduce time-to-value, shrink tuning burden, and credibly govern SaaS and AI workflows with automated remediation and auditable evidence.

## Sources and References


- *1 - AI adoption has reached 73% of enterprises in 2026, while real-time security governance is just beginning to emerge at 7%. (Netskope)*
- *2 - A commissioned study conducted by Forrester Consulting on behalf of Google, "Cloud Workers Are Key To Disruption Preparedness", 2020.*

Software Analyst®
Cyber Research

business personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

| | |
|---|---|
| CISO | VENDOR |
| SECURITY TEAMS | INVESTORS |