

THE AGENTIC DSPM SHIFT: HOW DATA DISCOVERY IS EVOLVING INTO REAL-TIME REMEDiation

LAWRENCE PINGREE





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 Us

Software Analyst Cyber 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.

Acknowledgements

Practitioners and CISOs,

We're excited to share our latest cybersecurity research. This report explores emerging technologies and market trends, offering practical insights to help organizations evaluate solutions, strengthen security, and make more informed decisions.

The Author

Lawrence Pingree is the Head of Research at SACR, where he leads research on data protection, AI security, identity, cloud security, and agentic security models. With over 16 years of industry analyst experience, including more than a decade at Gartner, he has authored over 300 research notes and provides research driven insights to enterprise security leaders on emerging cyber threats, cloud security, AI security, endpoint defense (EDR), and SD-WAN.

TABLE OF CONTENTS

Executive Summary	3
Key Findings	3
Agentic DSPM Report Thesis	4
Defining Emerging Agentic DSPM	7
Why Context Matters for Safe Remediation	7
The Great DLP Reset	8
From Visibility to Action, an Automation Maturity Model	9
What Agentic Should Mean in DSPM	10
Remediation-Oriented DSPM Architecture	12
Vendor Landscape: Varying Paths to Action	13
Agentic DSPM in Practice: Examining Remediation-First Workflows	15
Risks and Caveats of Agentic DSPM	19
Teleskope	22
Research Disclosure	25
Conclusion	26

Executive Summary

Data Security Posture Management (DSPM) excels at discovering, classifying, and mapping risks across cloud and SaaS data. However, as DSPM tools surface more findings, they create a major operational bottleneck where security teams are left manually investigating, routing tickets, and waiting on approvals to fix issues. As AI-driven workflows and modern collaboration environments increase the speed and scale at which sensitive data moves, these delays become more consequential.

The next phase of DSPM is reliable, governed action with AI and agentics. True Agentic DSPM goes beyond chatbot alerts to manage the full remediation workflow from context assembly and owner routing to change execution, proof, and rollback capabilities. Adoption will follow staged trust. Organizations can safely automate low-risk, reversible tasks like removing public links or redacting payment details, while reserving human approval and simulation for high-impact production changes. Ultimately, the future of DSPM is not just listing risks, but serving as a governed operating system for resolving them efficiently.

Key Findings

- DSPM visibility is necessary but no longer sufficient; remediation throughout is becoming a primary measure of value in the agentic enterprise.
- Agentic capabilities are only credible when they have the capability to reduce manual steps under defined authority and produce auditable outcomes.
- Reasoning and execution should be separated; probabilistic systems can interpret context, while deterministic controls perform approved changes.
- Telescope represents a remediation-first approach and is a key Native Agentic Remediation example in our analysis.
- Buyers should automate in stages, beginning with high-confidence, reversible actions and expanding authority only after operational evidence accumulates.

Agentic DSPM Report Thesis

DSPM is moving from a system that discovers and describes data risk to one that can safely resolve it through context-driven reasoning, governed action, and verified outcomes. The competitive boundary is increasingly defined by trustworthy, policy-controlled action across cloud, SaaS, collaboration and AI environments, not simply better classification, another conversational interface, or deterministic automation labeled as agentic.




Figure 1: Market Map – Data Security Posture Management (DSPM)

Why Now? DSPM Has a Remediation Gap

DSPM emerged because organizations lacked a dependable map of sensitive data across cloud object stores, data warehouses, databases, and SaaS applications. While DSPM brought vital discovery, classification, access mapping, and posture analysis, the category’s primary challenge still only begins once a finding is created.

An exposed dataset can be easy to identify but hard to resolve when security analysts must manually locate owners, evaluate downstream dependencies, obtain approvals, and navigate disparate administrative consoles. This visibility paradox means better discovery often increases backlogs, unless remediation capacity scales at the same rate.

Traditional separation between DSPM (focused on data at rest) and DLP (focused on data in motion) fails to address the end-to-end data-risk lifecycle. Modern buyers require continuous capabilities: understanding data, evaluating usage, and applying controls natively. Consequently, DSPM platforms must be evaluated not just by discovery coverage and classification accuracy, but by remediation depth, how natively actions are executed, how approvals or exceptions are handled, and what concrete evidence demonstrates a verifiable decrease in risk. The market is beginning to add Agentic DSPM operating models on top of continuous DSPM; these models overlap, and qualification should be assigned by production workflow rather than by vendor slogan.




Figure 2: The Three Eras of DSPM

Why Has the Data Environment Changed?

Sensitive information no longer stays mainly in databases and file repositories. It appears in Slack messages, Jira tickets, source code, meeting notes, design systems, support conversations, browser sessions and generated content. Collaboration systems are built for copying and sharing, which makes data useful, but also multiplies locations, permissions and exposure paths.

Critical Examples:

- AI intensifies this dynamic. A copilot can retrieve information from several repositories, summarize it into a new artifact and expose it to a broader audience.
- An agent can use delegated credentials, call tools and modify data across systems.
- Prompt and response logs can become new sensitive stores.
- Retrieval-augmented generation can surface data that was technically accessible but never intended for the requesting context.
- Every model connector, agent and MCP server introduces another route through which enterprise data can move.

External research confirms that data exposure is central to AI risk. In the [2025 CSA and Google Cloud survey](#), 52% of respondents identified sensitive-data exposure as their top AI-security concern.

[IBM found](#) that hybrid and multi-environment breaches cost more on average than on-premises-only breaches, reinforcing the impact of fragmented visibility and controls.

The [2026 Verizon DBIR](#) also highlights excessive cloud privileges and slow remediation of basic access weaknesses, showing that authorization hygiene remains unresolved even before agentic access is added.

Defining Emerging Agentic DSPM

Agentic DSPM goes beyond conversational AI, recommendation engines, and traditional automation by connecting contextual reasoning to governed remediation. A qualifying workflow must identify a risk, assemble the relevant data and business context, determine an appropriate response, apply approval and policy controls, execute the selected action, and verify the result. These capabilities must work together within a production workflow rather than exist as separate features across a vendor's portfolio.

Out of scope:

- Conversational interfaces that only summarize or surface findings
- Recommendation engines that do not reduce manual steps or govern execution
- Periodic discovery and posture reporting without remediation orchestration or execution
- Unbounded autonomous agents without scoped authority, deterministic controls, auditability, or rollback
- Standalone DLP enforcement lacking DSPM context, ownership, and posture intelligence

Why Context Matters for Safe Remediation

This environment raises the value of context. A classification label alone cannot determine whether an action is safe. The platform may need identity, ownership, business purpose, data lineage, destination, access path, requested operation and expected impact. Remediation timing also matters. A scheduled scan that identifies a public link tomorrow is valuable for posture management; it is not equivalent to detecting and removing the exposure when the sharing event occurs.

Periodic DSPM remains useful, especially for broad inventories and historical exposure. But the operating model must add continuous monitoring, event-driven assessment and response for fast-moving surfaces. The shift is from a static map of risk toward a feedback loop that understands, decides, acts and verifies.

SACR Key Takeaway: Automated remediation depends on more than knowing whether data is sensitive. Platforms must understand ownership, identity, intended use, exposure paths, and business impact before determining which action is appropriate.

The Great DLP Reset: How Modern DSPM and Agentic Remediation Drive Evolution Beyond Legacy Controls

Data Loss Prevention (DLP) is simultaneously undergoing a fundamental transformation in the great DLP reset, shifting from perimeter-centric blocking to context-aware enforcement.

Traditional DLP relied heavily on static regex rules and inline perimeter monitoring, resulting in high false-positive rates and operational friction for modern collaboration environments. In this evolved architecture, continuous discovery and classification serve as inputs into downstream action rather than standalone reporting metrics.

DSPM acts as the underlying truth layer for data security decisions, unifying data visibility with user identity and authorization context. Enforcement becomes distributed across native APIs, SaaS permission controls, real-time redaction pipelines, access modifications, workflow controls, and AI-layer interventions.




Figure 3: The Modern Data Security Architecture

From Visibility to Action, an Automation Maturity Model

[Maturity should be assigned by use case, not by platform slogan.](#) A product may be at Level 5 for redacting a payment-card number in Slack and at Level 2 for changing access to a complex business application. Likewise, a customer may deliberately cap a product at Level 3 until it has enough evidence to trust a specific remediation. Automation maturity describes how much authority a workflow has; Agentic DSPM qualification separately tests whether the workflow contains the required observation, context, reasoning, governance, execution, and verification loop. A mature deterministic automation workflow is not automatically Agentic DSPM.

The safest adoption path starts with deterministic, reversible and low-blast-radius actions. Simulation and observe-only modes should precede enforcement. Approval thresholds can then vary by confidence, sensitivity and operational impact.

The objective is not maximum autonomy; it is the least human involvement consistent with acceptable business risk. Below is a representation of the SACR Agentic DSPM Maturity framework.




Figure 4: Agentic DSPM Maturity Framework

What Agentic Should Mean in DSPM

Agentic systems are well suited to the investigative middle of the remediation workflow. A finding may require the platform to gather related resources, inspect access, read a policy, infer ownership and select an appropriate workflow. These are multi-step tasks that benefit from flexible reasoning and tool use.

Execution requires a different standard. If a platform removes access, deletes a file or changes a policy, the customer must know what changed, which identity authorized the action and whether the operation succeeded. A practical architecture separates a reasoning layer from an execution layer. The reasoning layer can interpret ambiguous context and recommend a response.

The execution layer uses controlled APIs, narrowly scoped permissions and deterministic runbooks to carry it out.

This distinction prevents agentics from becoming a synonym for unconstrained automation. The agent should not receive more authority than the workflow requires. High-impact actions should require stronger evidence, simulation or human approval. Every action should produce an audit trail with the input finding, policy basis, approver or service identity, target objects, result and rollback state.

Agentic Trust is the Gating Function

The largest barrier to automated remediation is not the absence of an agent. It is confidence in the underlying finding and in the proposed or automated change. A platform that misclassifies data, cannot identify the owner or lacks business context will not be trusted to revoke access automatically.

Classification quality, identity context and policy clarity are therefore prerequisites for autonomy.

Customers also need a preview of impact. Before broad enforcement, the product should support observe-only operation, simulation or campaign review. The user should see which objects will change, why they qualify, which dependencies may be affected and what happens if the action fails. Reversibility matters because even a correct security decision can create an unexpected business consequence.

Trust should be earned through operational evidence. For each action type, buyers should track recommendation acceptance, execution success, false-action rate, rollback frequency, mean time to verified closure and user disruption. A high-level statement that “90% of remediation is automated” is not meaningful without the denominator, action types and exception policy.

The most credible path is graduated authority. A program can begin with agent-assisted triage, advance to grouped campaigns, introduce approval-driven execution, and automate only the actions that repeatedly produce predictable outcomes. This is consistent with mature security-automation practice and more defensible than treating autonomy as a feature toggle.

Agentic DSPM qualification should be evaluated across six gates:

1. Observe: Event-driven or continuous observation in addition to scheduled inventory.
2. Understand: Assemble context across data sensitivity, identity, ownership, access, policy, destination, business purpose, lineage, AI usage, and operational impact as applicable.
3. Reason: Dynamically interpret the finding, select a specific remediation, and explain why that action is appropriate.
4. Govern: Apply bounded authority, approval thresholds, simulation or preview, exceptions, scoped credentials, and deterministic execution controls.
5. Act: Execute the selected remediation natively or through a tightly governed enforcement point; creating a ticket or launching a downstream workflow alone does not qualify.
6. Verify & Learn: Confirm target-state change, capture results and failures, preserve audit evidence, support rollback where available, and feed outcomes back into policy.

Remediation-Oriented DSPM Architecture

DSPM supplies the truth layer: what the data is, where it lives, who can reach it and why the exposure matters. DLP, SaaS APIs, cloud controls, identity systems and AI gateways provide enforcement points. The decision and governance layers connect intelligence to action without assuming that one product must own every control surface.

A remediation-oriented DSPM architecture can be represented as six connected layers:




Figure 5: Remediation-Oriented Representative Architecture

Vendor Landscape: Varying Paths to Action

The market should be compared by operating model, not by a single maturity rank. Vendors start from different strengths, DSPM, permissions analytics, DLP, governance or cloud security and may automate different portions of the lifecycle.

Vendor / Type	Starting Point	Remediation Question	Buyer Validation Priority
BigID	Emerging Agentic DSPM; data intelligence, governance and catalog foundation	Does broad context reduce remediation effort or remain primarily a discovery advantage?	Native actions versus downstream tickets; runtime versus after-the-fact controls.
Cyera	Emerging Agentic / Context-Driven DSPM; scaled DSPM expanding into DLP, identity and AI context	Does broader context produce faster remediation, or mainly improve discovery and prioritization?	Enforcement depth beyond supported public apps and Microsoft environments; connector coverage.
Teleskope	Native Agentic DSPM; remediation-first data security platform	Can contextual classification and specialist agents materially reduce work from finding to closure?	Actions available today by connector; approval model; rollback; failure handling; independent customer outcomes.
Varonis	Automated / Context-Driven DSPM; established permissions analytics and policy-driven access control	Which workflows are newly agentic versus mature permissions automation?	Coverage of modern SaaS/AI surfaces and execution evidence.
Veeam / Securiti AI	Emerging Agentic DSPM; DSPM + privacy + AI governance + resilience	Can graph-driven remediation and Veeam-backed recovery close risk without creating workflow friction?	GA vs preview Agent Commander; native vs downstream action; backup coverage/RPO; rollback proof; deployment burden.

Agentic DSPM in Practice: Examining Remediation-First Workflows

Telescope as a Representative Model of Remediation-First Agentic DSPM

To illustrate how the shift from data discovery to governed remediation is taking shape in practice, SACR examines Telescope as an example of a remediation-first data security platform. Its architecture, supported enforcement actions, and published customer deployments provide a useful basis for assessing how contextual understanding translates into operational outcomes. This section focuses on those workflows and their limitations, while the vendor profiles later in the report provide a standardized comparison across the broader market.

Telescope's approach centers on a Data Reasoning Layer that connects three stages: understanding data in business context, determining the appropriate response based on policy, and enforcing that response within the platform. Telescope is best assessed as a DSPM-led data security platform with native remediation and DLP workflows.

Its public positioning centers on a Data Reasoning Layer that connects three stages: understand the data in business context, decide the appropriate action from policy, and enforce the action in the same platform. [Available actions described by the company](#) include informing, redacting, quarantining, revoking access, relocating, encrypting and deleting.

The product's differentiation claim is not simply better discovery. It is that contextual classification and policy interpretation make automation safe enough to move beyond a ticket handoff. The company describes Prism for document-level understanding, more than 150 sensitive entity types, policy ingestion and native remediation across selected cloud, SaaS and AI environments. It also states that low-confidence cases route to human review and that actions are auditable and reversible.

Public customer evidence is strongest in bounded, high-confidence workflows. [Ramp reports](#) under-two-second detection in Slack, 100% automated redaction workflows and continuous historical visibility across Slack and Google Drive. [Telescope also publishes outcomes for The Atlantic](#), including a 95% reduction in time spent on data deletion and a 97% reduction in cost per deletion, and for EarnIn, including policy-driven self-service remediation. These references show the direction of value, but they remain vendor-hosted case studies. Buyers should request production demonstrations, reference calls and raw operational metrics.

SACR's independent research adds an important balance. [Prior evaluation in SACR'S April "The Great DLP Reset"](#) characterizes Telescope's architecture as API-driven and data-store-driven rather than a universal transparent inline proxy. That model can provide near-real-time response and native action in supported systems, but it should not be assumed to replace endpoint, browser, network or email DLP across every channel. Connector-specific permissions, latency and action depth must be verified.

Following a demonstration to SACR in September, [Telescope announced Kosmo's general availability on October 8, 2026](#). According to the company, Kosmo supports investigation, prioritization, remediation planning, policy development and reporting, but cannot directly modify customer systems without human approval. SACR therefore distinguishes the launch announcement from evidence of operational effectiveness. The relevant test is whether Kosmo reduces manual investigation, organizes related findings into actionable campaigns, helps establish accountable owners and proposes appropriate policies. Further test whether its proposed plans can be handed off to the platform for approved execution and verified closure, not merely whether it summarizes an alert.

The strongest buyer fit is an organization with sensitive-data sprawl across cloud and SaaS, a backlog of repetitive cleanup, and willingness to grant the connector permissions required for native action. Buyers seeking universal inline prevention, deep endpoint controls or model-security capabilities should assess complementary products.

Practical CISO Implementation Guide

Security leaders should treat Agentic DSPM as a governed operating-model change, not a feature toggle. The objective is to move selected data-risk workflows from visibility to verified closure while expanding automation only as evidence, context quality, and control maturity improve.

1. Baseline the remediation backlog and economics: Identify the highest-volume finding queues, current time-to-owner, time-to-action, exception rates, and the manual handoffs between detection and closure. Prioritize repetitive workflows where delay creates measurable exposure or analyst cost.
2. Establish the minimum trusted context: Before granting execution authority, validate classification quality, identity and ownership mapping, access paths, business purpose, policy state, destination, and operational impact. Weak context should cap the workflow at recommendation or approval-driven execution.
3. Map each workflow to an autonomy level: Classify use cases as Observe, Assist, Orchestrate, Governed Execution, or Selective Real-Time Enforcement. Assign maturity by workflow and connector rather than by platform-wide claims, and document which steps remain human-controlled.
4. Start with agent-assisted triage and simulation: Use agents to assemble context, group related findings, resolve likely ownership, recommend actions, and preview expected impact. Compare recommendations with analyst decisions before enabling automated changes.
5. Automate low-blast-radius, reversible actions first: Favor high-confidence actions such as removing public links, narrowing sharing, revoking stale access, or redacting clearly defined sensitive content. Require stronger approval for deletion, broad permission changes, or actions with material production impact.
6. Separate reasoning from execution: Allow probabilistic systems to interpret context and propose remediation, but use deterministic APIs, scoped credentials, policy checks, change windows, and explicit approval gates to perform the actual change.

7. Instrument the evidence loop: For every action, capture the input finding, policy basis, acting identity, approver where applicable, target object, execution result, failure state, rollback status, and confirmation that the intended target state was reached.
8. Measure trust before expanding authority: Track recommendation acceptance, execution success, false-action rate, rollback frequency, user disruption, mean time to verified closure, and exposure reduction. Expand autonomy only for workflows that repeatedly demonstrate predictable outcomes.
9. Validate connector-specific operational reality: Test native versus downstream actions, event latency, API permissions, rate limits, partial failures, rollback, and verification for the systems that matter most. Do not assume that API-based remediation provides universal inline coverage across endpoint, browser, network, email, SaaS, and AI surfaces.
10. Extend selectively into collaboration and AI workflows: Once the core governance and evidence loop is mature, apply the same controls to prompts, model inputs and outputs, retrieval sources, agent-accessible repositories, collaboration content, and AI-generated artifacts. Preserve human approval for novel or high-impact actions until operating evidence supports greater authority.

SACR Key Takeaway: CISO adoption should optimize for trustworthy risk reduction, not maximum autonomy. The practical end state is graduated authority: agents handle investigation and orchestration, deterministic controls execute bounded changes, and every workflow proves closure through auditable evidence before broader automation is granted.

DSPM Buyer Evaluation Checklist:

As DSPM products take on more remediation work, buyers need to look closely at what happens after a risk is identified. The important questions are practical: what the product can actually change, what context informs that action, where approval is required, and whether the platform can confirm that the issue was resolved.

The checklist below is designed to test those capabilities at the workflow level.

1. What can the product do natively, and what only creates a ticket or calls another platform?
2. Does the control run on a schedule, in response to an event, near real time or directly inline?
3. What data, identity, ownership and business context informs the decision?
4. Which actions require human approval, and can thresholds vary by risk and resource?
5. Can the customer simulate the policy and preview affected objects before execution?
6. Which identity performs the change, and how narrowly can its authority be scoped?
7. Are actions deterministic, auditable and reversible? What cannot be rolled back?
8. How are partial failures, rate limits, connector outages and conflicting policies handled?
9. Does the platform verify that the target state changed, or only report that a workflow launched?
10. What customer evidence shows reduced exposure, time saved or lower operational cost and what is the denominator?

Risks and Caveats of Agentic DSPM

- Agentic inflation: Vendors frequently apply the agentic label to superficial capabilities, such as basic conversational chat interfaces, alert summarization, semantic search, or passive recommendations. These interfaces merely surface existing findings without materially transforming the remediation workflow or executing governed actions to close risk.

- Automation error & blast radius: When automated agents are granted autonomous permissions to revoke access, quarantine assets, modify broad permissions, or delete files, a single false positive or contextual misinterpretation can escalate into a major operational outage. Evaluating business impact, user context, and operational dependencies is crucial before enabling automated execution.
- Incomplete enforcement coverage: API-driven and data-store-level remediation mechanisms offer rapid, native cleanup in supported cloud and SaaS platforms, but they operate out-of-band rather than directly inline. Consequently, API remediation alone does not provide universal inline prevention across endpoint, browser, web, network, email, or active data-in-use paths.
- Privilege concentration & non-human identity risk: Centralized remediation services requiring expansive admin rights across multiple cloud environments, data stores, and SaaS applications become high-value non-human identities (NHIs). If compromised or improperly scoped, these services expose the enterprise to significant blast-radius risks, necessitating strict identity governance and runtime security controls.
- Evidence quality & metrics validation: Vendor-published case studies, marketing benchmarks, and press releases serve as preliminary indicators of capability but cannot substitute for independently verified, production-level operational data. Security buyers must rigorously evaluate success rates, exception counts, total finding denominators, and mean-time-to-closure metrics within their specific technical environment.
- Category overlap & architectural alignment: Overlapping functionalities across DSPM, DLP, SSE, CNAPP, identity governance, and AI-security platforms increasingly target the same data-risk lifecycle. Rather than attempting to purchase a single category label, enterprise buyers should design a coherent, multi-layered architecture where DSPM functions as the contextual intelligence layer and enforcement is distributed across appropriate control points.

- Irreversible actions & recovery requirements: Destructive remediation activities, including permanent file deletion, destructive redaction, and sweeping entitlement revocations, carry far higher operational risks than reversible actions like public link removal.
- These high-impact operations require elevated approval thresholds, pre-execution simulation, robust backup/rollback mechanisms, and rigorous testing before deployment.




Figure 6: Market Map: Data Security Posture Management (DSPM)



teleskope

Vendor Overview

Telescope is a remediation-first agentic DSPM and data security platform that combines discovery, contextual classification, policy interpretation, and native action. Its differentiation in this report is not simply identifying sensitive data; it is using that context to move findings toward closure inside supported cloud, SaaS, collaboration, and AI environments. The platform is best suited to organizations with sensitive-data sprawl, repetitive cleanup work, and a willingness to grant connector permissions for direct remediation.

Use Case Performance

Telescope's Data Reasoning Layer connects business-context understanding, policy-based decisions, and remediation. The platform supports actions including informing, redacting, quarantining, revoking access, relocating, encrypting, and deleting. In SACR's September 2026 briefing, Telescope demonstrated Kosmo as an orchestrator agent with Search, Triage, Remediate, Reporting, and Policy Creation specialist agents; policy ingestion can generate policy suggestions from existing classification and retention documents. Customer examples discussed with SACR included redaction of PII from AI-generated meeting summaries in Google Drive, Slack redaction, and automated remediation of high-confidence findings.

Operational Reality

The architecture remains primarily API- and data-store-driven rather than a universal transparent inline proxy for Claude DLP, email, and retroactive audit. Telescope described Telescope-managed, single-tenant SaaS, and fully air-gapped deployment options. Telescope has published case studies of real-time Slack detection and redaction in a customer deployment, while broader capabilities include near-real-time Google Chat and Slack detection and deletion. In-path email DLP for Outlook and Gmail and broader data-lineage and identity coverage were described as roadmap areas in SACR's September briefing. Coverage, latency, approvals, rollback, and failure handling should therefore be validated by connector and action type.

Ticketing And Kosmo AI Agentic Remediation

Telescope supports ticket-based workflows; approval thresholds based on customizable per-customer policies and risk appetite; low-confidence review; how failures and partial changes are handled; and whether rollback is available for each action type. For Kosmo, Telescope offers each specialist-agent workflow and measures whether triage, policy creation, reporting, and remediation materially reduce manual work through verified closure.

Buyer Validation Priorities

Buyers should evaluate Telescope's remediation capabilities across their priority cloud, SaaS, and collaboration environments, focusing on which actions execute natively and which require human approval or additional workflows. Validation should cover connector-specific latency, configurable approval thresholds, exception handling, rollback, and evidence of successful remediation. Kosmo should be assessed for its ability to reduce manual investigation, triage, policy creation, and remediation planning, including how its recommendations translate into approved actions through Telescope's platform. Buyers should also request production examples demonstrating reduced manual effort, faster remediation, and verified closure.

SACR Perspective

Telescope represents a strong example of a remediation-first Agentic DSPM architecture in our current vendor sampling. Full qualification should be assigned by production workflow and connector, based on demonstrated contextual reasoning, governed execution, rollback where supported, and verified closure, not on the presence of an assistant alone. Organizations needing universal inline prevention across endpoint, browser, network, or email should evaluate complementary enforcement controls.

Research Disclosure

This report reflects SACR's independent research and analysis of the Data Security Posture Management (DSPM) market as of October 2026.

Teleskope participated in a sponsored research program associated with this report. Sponsorship provided access to vendor briefings, product demonstrations, and supporting materials.

SACR's research also incorporated enterprise customer validation, proprietary market telemetry, and broader market analysis. Sponsorship did not influence SACR's conclusions, analysis, or editorial direction.

The vendor profiles in this report are not product reviews, rankings, or procurement recommendations. They are intended to illustrate how different vendors are responding to the market developments and challenges discussed throughout the report. All opinions expressed are those of SACR.

Conclusion

Data discovery and classification remain foundational, but market differentiation is increasingly shifting toward remediation throughout and governed action.

The future of DSPM relies on combining probabilistic reasoning over context with deterministic enforcement controls, ensuring security operations keep pace with AI-driven data proliferation.

Some vendors are beginning to use agents to assemble context, investigate findings, select actions, and coordinate remediation, while DLP, SSE, cloud security, and governance platforms approach the same challenge from adjacent control points.

Agentic DSPM qualification should therefore be assigned by production workflow, not inferred from AI features, deterministic automation, or portfolio breadth.

SACR Key Takeaway

The winning platforms will not be those with the most detections or the most fluent assistants, but those that prove success, connect trustworthy data intelligence to governed action, and allow customers to expand autonomy through evidence.

For buyers, the central question is no longer just what a platform can find, but what it can safely resolve, under whose authority, and how risk reduction will be measured.



business



personal



softwareanalyst.io



Trusted research. Sharp insights. Real conversation.

CISO

VENDOR

SECURITY
TEAMS

INVESTORS