

THE NEW DATA SECURITY MATURITY MODEL FOR THE AI ERA

| 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 a new data security maturity model for the AI era. This report provides a practical path for connecting data visibility, identity context, enforcement, and governed automation.

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

Actionable Summary	3
SACR Data Security Maturity Model	4
Why Data Security Maturity Has Stalled	8
Invisible Ceiling: Why Maturity Stalled	12
Cyera's Operational Maturity Roadmap	19
CISO Capability Assessment Framework	23
FAQs	25
Best Fit Customer Profiles	29
SACR Takeaway/Methodologies	32
Practical Recommendations for CISOs	33
CISO Maturity Checklist	33

Actionable Summary

Data-security maturity is not defined by the tools an organization owns. It is defined by what the organization can repeatedly prove: where sensitive data resides, who or what can access it, which exposures matter, how remediation is governed, and whether risk reduction can be validated.

Key takeaways

- Mature programs connect discovery, classification, identity and access context, enforcement, remediation, governance, and measurement into a continuous control loop.
- DSPM provides the contextual truth layer, while DLP and adjacent controls provide enforcement paths; neither produces mature outcomes in isolation.
- Progress depends on people, process, information, and technology, not technology deployment alone.
- Agentic systems can reduce the cost and operating effort associated with advanced maturity, but automated action must remain policy-constrained, explainable, auditable, reversible where appropriate, and continuously validated.
- In the AI era, the strategic objective expands from protecting stored data to governing what users, applications, non-human identities, and agents can see and do.
- This report defines a five-level maturity path, explains why programs stall, describes the architecture needed to connect context with enforcement, and provides practical measures CISOs can use to evaluate progress.

SACR Thesis

SACR defines maturity not by tool adoption alone, but by what an organization can prove, where sensitive data lives, who or what can access it, which exposures matter most, how remediation is governed, and whether outcomes can be validated over time. Achieving true data security maturity requires a fundamental operating-model shift rather than a piecemeal assembly of analyst-defined product categories. Advanced maturity demands that discovery, classification, protection, detection, remediation, and governance operate as a connected, unified control loop. Current programs often stall at Level 3 due to retrospective scanning, brittle policies, and operational drag.

Core Challenges Faced by Enterprises

- **Distributed Data and Complex Actors:** Sensitive data sprawl across SaaS, cloud, and AI environments combined with a proliferation of human and non-human actors renders legacy, perimeter-focused security models insufficient.
- **Architectural Silos vs. Unified Control:** Programs often stall because discovery (DSPM) and enforcement (DLP) function as context-poor, siloed tools; true maturity requires integrating these into a continuous control loop of discovery, protection, remediation, and [governance](#).
- **Bounded Agentic Governance:** The rapid adoption of AI and agentic workflows introduces novel risks such as excessive agency and prompt injection, necessitating that security automation be strictly policy-constrained, explainable, auditable, and human-in-the-loop.
- **Provability as the Maturity Metric:** True data security maturity is not defined by tool ownership, but by the organization's demonstrable ability to prove data location, verify access, and validate that remediation governance effectively reduces risk.

SACR Data Security Maturity Model




Figure 1: SACR Data Security Posture Maturity

The SACR Data Security Maturity Model outlines how an organization's data protection capabilities evolve through five distinct stages. By benchmarking current practices from ad hoc, manual interventions at Level 1 to AI-driven, predictive security at Level 5, this framework helps security teams assess their overall posture, identify critical limitations, and chart a strategic roadmap toward automated, continuous compliance. Our Maturity Model describes an organization's progress across four interdependent dimensions:

- **Technology:** The coverage, integration, precision, and control capabilities available.
- **People:** Defined accountability across security, IT, data owners, application owners, and executive leadership.
- **Process:** Repeatable workflows for prioritization, enforcement, remediation, exceptions, escalation, and validation.
- **Information:** Reliable measures that demonstrate coverage, decision quality, operating performance, and risk reduction.

Here, maturity refers to an organization's demonstrable ability to manage data risk across its entire lifecycle. It is not merely about deploying security tools, but about the organization's capacity to prove that it knows where sensitive data resides, who has access to it, how it is being used, how risk is being prioritized, how remediation is being governed, and how these outcomes are being validated over time.

Note: An organization has achieved a maturity level only when it can produce evidence across all four dimensions. Technology deployment alone does not establish maturity.

Level 1: Ad Hoc Visibility (Initial)

- **Core Problem:** The practitioner needs to solve for basic, uncoordinated, reactive visibility across static and fragmented environments.
- **Proof Required:** Demonstrate the capability to identify where sensitive data lives and verify basic data inventory accuracy.
- **Capabilities & Operating Changes Needed:** Transition from manual spreadsheet tracking to initial automated discovery tooling.
- **Platform Relevance:** Provides a significantly more accurate and automated inventory than manual spreadsheets to establish a trustworthy baseline truth layer.

Level 2: Baseline Discovery (Repeatable)

- **Core Problem:** The practitioner needs to solve for consistent, repeatable process coverage across core data stores.
- **Proof Required:** Prove that discovery and classification processes run consistently across core repositories on a scheduled or continuous basis.
- **Capabilities & Operating Changes Needed:** Standardize data classification policies and centralize audit logging across storage layers.
- **Platform Relevance:** Automates continuous discovery and classification to enable scalable, consistent policy application without heavy manual overhead.

Level 3: Risk-Driven Policy Alignment (Defined)

- **Core Problem:** The practitioner needs to solve for policy noise, alert fatigue, and a lack of risk prioritization.
- **Proof Required:** Demonstrate risk-based prioritization and clear alignment between policy definitions and actual discovered risk.
- **Capabilities & Operating Changes Needed:** Shift from isolated tool silos to integrated discovery and real-time policy enforcement mechanisms.
- **Platform Relevance:** Correlates deep data context (sensitivity, identity, and access entitlements) to improve enforcement precision and drastically reduce false positives.

Level 4: Operationalized Remediation (Managed)

- **Core Problem:** The practitioner needs to solve for operational drag, manual triage bottlenecks, and uncoordinated remediation workflows.
- **Proof Required:** Prove measurable reductions in sensitive-data exposure alongside validated, SLA-tracked remediation cycle times.
- **Capabilities & Operating Changes Needed:** Integrate findings directly into enterprise workflow engines, establish explicit ownership, define SLAs, and operationalize exception management.
- **Platform Relevance:** Bridges the gap between posture insights and execution by facilitating closed-loop remediation workflows and automated owner assignments.

Level 5: Governed Adaptive Data Security (Optimized)

- **Core Problem:** The practitioner needs to solve for autonomous, safe scaling of data security without human operational bottlenecks or delayed intervention.
- **Proof Required:** Demonstrate that automated decisions and actions are policy-constrained, explainable, auditable, reversible where appropriate, and continuously validated. Evidence should show what sensitive data agents can access, which identity and purpose context informed each decision, what action was taken, and whether the result reduced risk without unacceptable business disruption.
- **Capabilities & Operating Changes Needed:** Adopt agentic, policy-constrained orchestration with robust audit boundaries, clear approval gates, and rollback mechanisms.
- **Platform Relevance:** A data-security platform can support this level by combining data, identity, entitlement, purpose, and policy context to govern what users, applications, non-human identities, and agents can see and do.




Figure 2: Data Security Maturity Model

The model should be treated as an operating-model progression, not a product checklist. It is consistent with the broader idea that data security maturity should develop through staged improvements in governance, protection, detection, response, and recovery.

A Historical Perspective: Why Data Security Maturity Has Stalled

The history of enterprise data security is defined by architectural misalignment: sensitive data is becoming distributed, dynamic, and cloud-mediated faster than controls could adapt. As workloads moved across cloud, SaaS, collaboration platforms, developer environments, endpoints, and now AI-enabled workflows, many data-security programs remained anchored to static policies, fragmented inventories, siloed toolsets, and control points that lacked enough context.

The result is not that DLP or DSPM failed. The better argument is that [each is incomplete in isolation](#). DLP cannot be effective in isolation because accurate enforcement increasingly depends on sensitivity, ownership, identity, entitlement, business-purpose, and workflow context. DSPM can discover and prioritize risk, but it must connect to enforcement, workflow, and validation mechanisms to change outcomes. Fragmented DSPM and DLP programs result in false positive noise and struggle to drive action. Nonprofit and [standards-based guidance](#) increasingly points to the same operating-model requirement, while security programs need connected governance, protection, detection, response, and recovery processes rather than isolated tools. To transcend basic controls and achieve an optimized Level 5 posture, organizations must move away from brittle, reactive tools. Leveraging semantic classification, dynamic context, and agentic automation provides the foundation for continuous threat prevention operating at modern enterprise speed.

The Current State of DLP Technology (2026)

Many organizations appear to sit between baseline discovery and early risk-driven prioritization. They may have DLP coverage across email, endpoints, SaaS, or network channels, and they may also have discovery or classification tooling. But they often struggle to connect discovery, access context, identity context, policy enforcement, workflow ownership, and remediation validation into a single operating model. Rules that rely on regex, along with generic LLMs, may detect some regulated business data types, but miss proprietary data and generate high false positives.

DSPM Foundation: The Truth Layer for Agentic Security

While governed automation may eventually help organizations scale remediation, its effectiveness depends on a robust [Data Security Posture Management \(DSPM\) foundation](#). DSPM serves as the data-security truth layer, the data inventory and context needed to understand where sensitive data lives, what it is, who or what can access it, how it is exposed, and why it matters.

This foundation enables the creation of unified data, action and context graphs, which help correlate sensitive data with both user and Non-Human Identity (NHI) behaviors and any associated agentic actions in the case of agentic identity.

These merged context graphs can provide the missing ingredients for better security decisions: sensitivity, ownership, access, exposure, identity behavior, policy, business context, and remediation history.

The goal should not be an unconstrained immune system, but a governed data-security control loop where discovery informs prioritization, enforcement, remediation, and validation. It paints a picture of how data is actually being used, whether that aligns with business intent, how it might be exposed, why the data matters, and what the risk or impact is.

Without this visibility and context, even the most advanced enforcement policies remain a form of architectural guesswork (often producing extensive false positives), forcing security teams to apply rules to data they haven't fully discovered or verified.

In prior research, SACR articulated the need for a truth layer in data security technology deployments, offering buyers both an intelligence and enforcement layer for DLP and DSPM combined ([The Great DLP Reset: Securing Data in the Age of SaaS, Cloud, and AI](#)).




Figure 3: The DLP Control Plane (Reference: The Great DLP Reset)

By integrating DSPM as an essential prerequisite, organizations help empower business use cases with greater data context. It helps data security platforms transition from broad, brittle rules to surgical precision and data centricity, foundational for agentic and autonomous operation.

This foundational visibility feeds the moment of truth at data security product ingestion, ensuring that, when agentic-enabled solutions for DSPM and DLP evaluate a data access request or an upload event, they do so with a complete understanding of the data's lifecycle, lineage, and overexposure risks. Ultimately, DSPM provides the map.

DLP, SaaS controls, endpoint controls, identity systems, workflow tools, and AI-security controls provide enforcement and remediation paths. The future opportunity is a unified architecture where discovery continuously informs prevention, prioritization, and validated remediation.

CISO implication: Understanding the DSPM foundation empowers CISOs to treat data discovery as an indispensable truth layer for all automated security actions. CISOs should evaluate whether their architecture can reliably connect data sensitivity, identity, entitlement, exposure, business purpose, policy, and remediation history, directing their engineering teams to measure the exact time gap between sensitivity discovery and active enforcement.

Outcome metrics buyers should track:

- Time from sensitive-data exposure creation to detection
- Time from detection to owner assignment
- Time from owner assignment to validated remediation
- Number of unresolved high-risk sensitive-data exposures
- Rate of excessive access to sensitive repositories
- DLP false-positive and false-negative trends by control point
- Percentage of exceptions reviewed, approved, and expired on schedule
- Percentage of remediations validated after execution

The Gap: Traditional systems are reactive because they wait for data exfiltration attempts to manifest before responding, and they do not understand the data or access context before movement happens.

Buyers often struggle to:

- Adapt policies safely without approval boundaries and audit trails
- Correlate sensitivity, identity, access, exposure, and business context quickly enough to guide action
- Convert findings into governed remediation with ownership, SLAs, exceptions, and validation
- Detect novel data-exposure patterns, AI-workflow misuse, excessive agency, and policy-bypass attempts
- AI and agentic workflows make this harder because they create new pathways for sensitive-information disclosure, prompt injection, insecure tool use, excessive agency, and uncontrolled data access.

The Invisible Ceiling: Why Maturity Stalled

Why DSPM and DLP Are Incomplete as Separate Operating Models

DSPM establishes data context. DLP and adjacent controls provide enforcement. Identity, workflow, governance, and validation connect the two into measurable risk reduction.

The Trust Layer: Governing What Agents Can See and Do

In AI-enabled environments, data-security decisions must account not only for data sensitivity, but also for the identity, purpose, permissions, tools, and intended actions of the agent requesting access. The trust layer is the decision system that combines this context to determine what an agent may discover, retrieve, transform, or transmit and under which conditions (also, [see SACR research on ARISE](#)). A mature trust layer does not merely produce a risk score. It supplies evidence that the resulting decision was appropriately contextualized, policy-constrained, explainable, and effective.

CISOs should evaluate the quality of this decision layer by asking:

1. What sensitive data can an agent discover, retrieve, transform, or transmit?
2. Which identity, entitlement, purpose, and policy context informed the decision?
3. Was the requested action allowed, warned, blocked, escalated, or remediated?
4. Can the organization explain and audit the decision?
5. Can the action be reversed where appropriate, and can its outcome be validated?
6. How are false approvals, false denials, exceptions, overrides, and downstream outcomes measured?
7. Does feedback improve decision quality without allowing policy to change outside approved governance boundaries?

The Identity Substrate: Governing Non-Human Actors

The effectiveness of an Agentic solution relies on a robust identity substrate focused on Non-Human Identities (NHI) and service-account behavior.

In modern cloud and AI-enabled environments, non-human identities including service accounts, workloads, API keys, automation, and agents can accumulate broad or poorly governed access to sensitive data.

Governing these actors requires moving beyond static credentials to intent-aware, ephemeral access.

By baselining the behavior of service accounts and machine identities, Agentic systems can detect anomalies, such as an agent attempting to aggregate PII across disparate systems, and intervene at machine speed, ensuring that autonomous actors remain constrained by business-aligned risk logic.

Agentic workflows also introduce risks such as [excessive agency, insecure tool use, prompt injection, and sensitive-information disclosure](#), which is why [identity governance and data governance need to be evaluated together](#).

The Result: a staged path toward governed automation

Some organizations may be able to automate bounded remediation workflows over a multi-quarter roadmap, but Level 5 maturity depends on data coverage, ownership quality, integration depth, workflow discipline, governance readiness, auditability, and tolerance for automated action.

Data Security Program Goal Examples:

- **MTTD:** <1 hour (vs. industry average of 24+ hours)
- **Threat detection rate:** >90% (vs. <30% for signature-based systems)
- **False positive rate:** <1% (vs. 10–25% baseline)
- **Reduced unresolved high-risk data exposure:**
 - Faster owner assignment for sensitive-data findings
 - Shorter remediation cycle time
 - Lower DLP policy noise in specific channels
 - Better exception governance
 - Higher percentage of validated remediations
 - Clearer approval boundaries for automated actions
 - Better visibility into AI and non-human identity access to sensitive data

Agentic Data Security Architectures Are a Core Differentiator

Decoupled Agentic Platform Capability Stack

The more defensible architecture separates core functions into interoperable services: discovery, classification, identity context, risk assessment, policy decisioning, enforcement, workflow, audit, and validation.

AI may improve parts of this system, but the maturity goal should be governed automation rather than unrestricted autonomy. [Credential abuse appears in 39% of breach progressions, even though it was 13% as the initial access vector.](#)

Credential abuse and identity compromise remain important data-security concerns because sensitive-data exposure often depends on who or what can access the data, not just where the data resides. Remediation efforts are still a critical aspect of moving these numbers down and reducing risks, especially as agents proliferate under less mature governance than users have historically had to date, meaning visibility alone is not enough, and we must take remediation action. The operating model should support human review where risk is high, automated recommendation where confidence is moderate, and automated execution only where actions are low-risk, reversible where appropriate, and auditable.

Level 5: How the Capability Stack Supports Governed Automation




Figure 4: Agentic Architectures Are a New Core Differentiator.

It combines:

- **Discovery (Ingestion):** Operates at the moment of entry, scanning content immediately during upload (e.g., to AWS S3, Azure Blob, or GCS) to enforce real-time policies before data becomes broadly accessible, eliminating the inherent risks and security lag of retrospective scanning.

- **Classification (AI-native):** Leverages NLP and generative AI models to achieve deep semantic understanding, accurately interpreting complex relationships between data elements (e.g., recognizing "Customer Name" + "SSN" as PII) and achieving >95% accuracy where standard regex and static patterns fail. Uses a mix of patterns, labels, metadata, semantic analysis, and contextual signals to identify sensitive data and improve policy decisions. Any accuracy claim should be validated against the customer's own data, with precision, recall, false-positive, and false-negative rates measured transparently.
- **Risk Assessment (AI-enabled):** Dynamically computes entropic risk scores in real time by continuously analyzing user and Non-Human Identity (NHI) behavioral anomalies, network topology changes, threat evolution, and regulatory imperatives (e.g., GDPR, HIPAA) to execute automated or human-in-the-loop remediation. Combines data sensitivity, exposure, access, identity behavior, business context, regulatory obligations, and workflow history to recommend whether an action should be allowed, warned, blocked, escalated, or remediated. Automation should remain policy-constrained and auditable.

One Key Innovation Is Reducing the Gap Between Exposure and Action

To move beyond reactive, signature-based security, organizations must adopt a fundamentally new approach to data defense. The following section details the specific architectural components of the agentic-enabled data security platform that enable this shift, moving from retrospective, scan-based detection to real-time, ingestion-time prevention.

By decoupling core services and applying context-aware intelligence, this approach can reduce the time between data exposure and policy-aware action. The goal is not to claim that all threats can be stopped at entry, but to make data-security decisions earlier, more contextual, and easier to validate. Intelligent data security platforms that are AI-native and agentic are emerging to take advantage of the latest generation of techniques for applying policy to context.

Intelligent Data Security Platforms (AI- and Agentic-Enabled)




Figure 5: Intelligent Data Security Platform

The Challenge with Traditional DLP

- **Delayed Context:** Some sensitive data is not classified, governed, or remediated until after it has already been created, moved, shared, or exposed.
- **Reactive Posture:** Breaches are often identified weeks later through historical log analysis rather than real-time prevention. Security teams may discover exposure through alerts, tickets, audits, or historical analysis rather than through a continuous operating model that connects discovery to action.
- **Operational Drag:** Retraining monolithic ML models to address new threats typically requires 2–4 weeks of system downtime. Policy tuning, exception handling, alert review, and remediation coordination can consume significant security and business-owner time when controls lack context or workflow integration.
- **Limited Context:** Traditional DLP often lacks the context to distinguish a low-risk event from a business-critical exposure, which is why teams end up with excessive noise or overly blunt blocking policies.

The Governed Automation Advantage: Earlier, More Contextual Decisions




Figure 6: Real-Time Data Security Flow

The Result

Data should be classified, evaluated, and governed as close as practical to the point of creation, movement, access, or use. This reduces reliance on after-the-fact discovery alone and gives enforcement systems more context before they warn, block, quarantine, escalate, or remediate.

The advantage is not AI-native branding by itself. The advantage is better context, sensitivity, identity, access, business purpose, exposure, workflow history, and policy intent.

The maturity model above should stand independent of any single vendor. Cyera is useful as an example because it shows how DSPM-derived context can be positioned to improve DLP precision, AI-data governance, identity-aware risk prioritization, and remediation workflows.

Cyera and Data Security Maturity: A Strategic Roadmap for Operational Maturity

Cyera Platform Summary

While the market may still perceive Cyera primarily through the lens of the DSPM category, the platform's architectural scope extends across broader data security, modernized DLP, identity, agent context, and automated remediation. Rather than being confined to data posture discovery alone, Cyera operates as an integrated control layer designed to align sensitivity, identity, access, and remediation workflows.

The company's strategic direction extends directly toward controlling what internal and external AI agents can see and do, grounding data access in verifiable policy. This operational scope directly fulfills the trust layer and decision-quality practitioners need, ensuring that every data security decision is informed by deep context before enforcement or remediation occurs.

CISO implication: Evaluating vendor examples like Cyera assists CISOs in selecting platform partners that actively connect data context to modernized DLP and identity controls. CISOs should direct procurement to validate vendor accuracy and workflow capabilities directly against internal enterprise data.

Cyera Core Insight

The data security market is moving through a DLP reset and a program reset focused on addressing AI agent risks and threats. For years, DLP programs tried to protect sensitive data through static policies, pattern matching, manual tuning, and enforcement decisions made without enough understanding of the data itself.

That model breaks down in cloud and AI-era environments because sensitive data is no longer concentrated in a small number of managed repositories or predictable movement paths, and the context of that data increasingly affects both risk and prioritization.

The new control plane has to answer a different question:

What sensitive data matters, who or what can reach it, how is it exposed, what business process depends on it, and what action will reduce risk without breaking the business?

That question sits between DSPM, DLP, CNAPP, identity security, data governance, and AI security. No adjacent category owns the whole problem. CNAPP understands cloud posture. Identity security understands users, privileges, and entitlement paths. DLP understands enforcement points. DSPM understands the data estate, while agentic visibility and control extend this context and control to agents. The maturity opportunity for practitioners is the control layer that connects these signals towards real-time actionable decisions.

The maturity path customers are trying to travel takes them:

- From incomplete visibility to reliable sensitive-data discovery.
- From classification to data-risk prioritization.
- From DLP noise to policy precision.
- From tickets to closed-loop remediation.
- From manual workflows to governed automation.
- From uncertainty about what AI can access to governed AI adoption.

Market Context:

Why legacy DLP is being reset

The cost of data breaches continues to rise, growing 12% year over year up to \$4.99M average across industries, and AI-driven attacks increased 56% over [last year's study](#). Data breach costs and AI-related security risks continue to put pressure on data-security programs. Legacy DLP was built for a more bounded enterprise. It assumed that sensitive data could be classified through known patterns, monitored through a limited number of channels, and governed through policies that security teams could manually tune over time. That assumption is under pressure.

Data now spreads across cloud storage, SaaS applications, databases, collaboration platforms, developer environments, AI tools, and third-party workflows. Users, service accounts, applications, and agents can all interact with sensitive data.

The risk is no longer only “Did someone exfiltrate a file?” It is also:

- Is sensitive data overexposed?
- Is access excessive?
- Is the wrong identity touching the wrong data?
- Is data being used in an AI workflow without appropriate controls?
- Is DLP blocking the wrong things and missing the right ones?
- Can the organization prove that remediation reduced risk?

This creates the market opening for DSPM-led data security platforms. [Microsoft reports that 32% of surveyed organizations’ data security incidents involve use of AI tools. It also says 47% of surveyed security leaders are implementing generative AI-specific controls, up 8% since the 2025 report.](#)

The promise is not merely broader discovery. AI and agentic workflows expand the data-security attack surface by introducing new risks such as sensitive-information disclosure, prompt injection, insecure tool use, excessive agency, and uncontrolled access to enterprise data.

Those risks strengthen the case for data context, identity context, policy enforcement, and governed automation before AI-enabled workflows are broadly trusted.

The promise is better decision-making. Discovery only matters if it changes prioritization, enforcement, remediation, auditability, or executive understanding of risk.

Vendor Thesis: Where Cyera fits

Cyera is positioned to compete for the data-security truth layer: the system that helps customers understand sensitive data, exposure, access, AI usage, and risk context across modern environments. Its opportunity is strongest when that context improves DLP precision, access governance, remediation prioritization, and evidence of risk reduction for modern DLP, AI agent and data security programs.

Although integrations must still occur with the Oasis acquisition, it places Cyera in an advantaged position to own unified protection of data and agentic identities, interactions, access control and data lineage across users, NHIs including AI chat and AI agents. The [Oasis acquisition](#) could strengthen [Cyera's position in AI and agentic identity security](#), but buyers should separate current capability from roadmap integration. Buyers should validate which capabilities are generally available today, which are integrated into the platform, and which remain roadmap or services-led.

SACR Analyst assessment: The key is that Cyera's strongest narrative is to make DLP and data security programs more precise, contextual, and operational, establishing a data foundation required to safely adopt AI. Existing DLP controls may still perform the block, warn, quarantine, encrypt, or monitor action. Cyera's value is higher if it improves the quality of those decisions and helps customers prove risk reduction.

Analysis:

Data Security Maturity Levels by Benefit

Maturity Levels must consider multiple dimensions across:

- Technology capability
- People and ownership
- Process and workflow
- Information and measurement
- Evidence that the level has actually been achieved

Level	Capability	Gaps and Challenges
Level 1: Initial – Ad hoc visibility	Automated discovery tooling providing an accurate data inventory baseline across core stores.	Manual spreadsheets and reactive response mechanisms indicate uncoordinated visibility.
Level 2: Repeatable – Baseline Discovery	Continuous discovery, standardized classification policies, and centralized audit logging.	Lack of real-time policy enforcement and reliance on batch and scheduled scans.
Level 3: Defined – Risk-Driven Alignment	Correlates data sensitivity, identity, and access entitlements to reduce alert noise.	Remediation workflows remain fragmented across non-security stakeholders.
Level 4: Managed – Operationalized Remediation	Closed-loop remediation workflows, automated owner assignments, and SLA tracking.	Requires proof of cross-functional adoption, SLA performance, and post-remediation validation.
Level 5: Optimized – Governed Automation	Policy-constrained orchestration, trust layer controls, and continuous posture optimization.	Requires evidence of safe, autonomous action with auditable, reversible rollback controls.

CISO Capability Assessment Framework

Buyers should evaluate their DSPM and DLP solution alignment and maturity programs by capability dimension, evidence strength, and maturity enablement. The goal is to help a CISO move from one maturity stage to the next.

CISOs should ask vendors to prove:

1. Coverage: Which cloud, SaaS, database, endpoint, collaboration, AI, and on-prem systems are actually covered?
2. Classification quality: What are the measured precision and recall rates against the customer's own data?
3. Identity context: Can the platform correlate sensitive data with users, service accounts, non-human identities, agents, entitlements, and actual usage?
4. DLP decision quality: How does data context improve enforcement decisions, reduce noise, or prevent unnecessary business disruption?
5. Remediation workflow: Can findings be assigned to owners, tracked through SLAs, exception handling with governance, and validated after completion?
6. Automation boundaries: Which actions can be automated, which require approval, and how are actions logged, reversed, or validated?
7. AI workflow governance: Can the platform detect sensitive-data use in AI tools, prompts, agents, retrieval systems, and model-connected workflows?
8. Evidence of risk reduction: Can the customer prove reductions in unresolved exposure, excessive access, remediation cycle time, policy noise, and exception backlog?

Data Security Evaluation Framework




Figure 7: Data Security Evaluation Framework

FAQs: Where do leading platforms have differentiated depth versus adjacent DSPM, DLP, CNAPP, and identity security vendors?

Leading platforms differentiate when they provide data-native context that adjacent categories do not collect or do not use deeply enough. The defensible differentiation is not generic cloud risk, generic identity risk, or generic policy enforcement.

It is the ability to answer harder operational questions like:

1. What sensitive data exists?
2. Why does it matter?
3. Who can reach it, who actually has touched it, and does that match intent and appropriate use?
4. How is it exposed?
5. What policy applies?
6. What action will reduce risk with the least disruption?
7. What other events surround the moment of transaction, but can refine decisions?
8. Can I automate and enforce data loss prevention actions with humans in the loop and without?
9. Can I control access to data from internal and external AI agents and service providers?
10. Can I kill-switch a malfunctioning AI agent or stop an exfiltration attempt by an AI agent or a user?

Each adjacent category contributes a different part of the control plane. The emerging platform opportunity sits between these categories. The leading systems will be those that correlate, manage risk and control and prevent sensitive data leakage or data with excessive access, external sharing, weak permissions, privileged identities, cloud exposure, regulatory obligations, and active movement. Platforms that only produce another inventory or alert queue will be less differentiated because they do not change the decision or accelerate the desired outcome.

- **DSPM:** Is strongest when the decision problem is discovery, classification, ownership, exposure, and risk prioritization across data stores.
- **DLP:** Is strongest when the decision problem is enforcement at the point of movement, sharing, transfer, or attempted exfiltration.
- **CNAPP:** Is strongest when the decision problem is cloud posture, workload exposure, entitlement paths, vulnerability context, and runtime infrastructure risk.
- **Identity security:** Is strongest when the decision problem is users, privileges, entitlements, access reviews, governance, and identity threat paths.

Which customer maturity stages are best served by current market capabilities?

Current market capabilities are strongest at **Level 2** and **Level 3**.

Most modern DSPM platforms can move customers from ad hoc visibility to baseline discovery. They can identify sensitive data across major environments, classify it, and expose obvious concentrations of risk. Leading platforms can also help customers move into risk-driven prioritization by combining sensitivity, exposure, identity and access, sharing, and compliance signals.

Signals: The market is weaker at Level 4. Many vendors can create tickets, assign owners, and recommend actions. Fewer can prove that remediation was completed, validated, measured, and repeated across teams. That gap matters because operationalized remediation is the point where posture management becomes risk reduction.

The observed market weakness at Level 4 stems from a dual challenge involving technology limitations and organizational maturity. While some vendors provide the necessary ticket-creation capabilities, they often fail to integrate seamlessly into complex cross-functional workflows, leaving security, IT, and compliance teams to manually bridge the gaps between discovery and remediation.

Organizations themselves frequently struggle because they continue to operate in silos, where security teams own discovery, but business units own the data, creating friction in owner assignment and accountability. Many organizations have not yet shifted their metrics from effort-based such as the number of tickets closed to outcome-based measurements like validated risk reduction over time.

Operationalizing Level 4 requires vendors to offer better workflow orchestration while simultaneously requiring organizations to break down internal silos and adopt governance models that prioritize measurable risk outcomes over simple activity metrics.

Level 5 remains early. Agentic remediation is strategically important, but most organizations will require approval gates, workflow controls, rollback paths, and strong evidence before trusting autonomous changes to sensitive data, access, or enforcement policy.

What capabilities are required to move from visibility to operationalized remediation?

Moving from visibility to operationalized remediation requires a change in operating model. Findings must become governed work. Owners must be accountable. Outcomes must be validated. Evidence must be preserved. The maturity shift is simple: from finding sensitive data to governing sensitive-data decisions.

Required capabilities include:

- Risk-based prioritization.
- Ownership mapping.
- Actionable remediation guidance.
- Workflow integration.
- SLA and exception management.
- Closed-loop validation.
- Evidence generation.
- Metrics and executive reporting.

How should agentic remediation be evaluated in a data security maturity context?

Agentic remediation should be evaluated as governed execution, not as an AI feature. The relevant question is not whether an agent can suggest a fix. It is whether the agent can safely recommend, coordinate, or execute a data-risk reduction action while preserving accountability.

Outcome: In this model, a chatbot that suggests fixes is Level 3 support. An agent that orchestrates governed remediation and proves risk reduction can support Level 4. An agent that continuously optimizes posture inside policy, approval, rollback, and audit boundaries begins to support Level 5.

How should DLP modernization be scored when DSPM provides richer data context but enforcement still occurs through existing DLP controls?

DLP modernization should be scored by whether DSPM context improves the precision, relevance, and outcomes of enforcement. The DSPM platform does not have to replace the DLP enforcement point. It has to make enforcement better.

Outcome: The high score belongs to the platform that makes DLP more targeted, adaptive, risk-aware, and measurable, even if the block, warn, or quarantine action still occurs in another control.

What customer evidence would prove movement from Level 2 or Level 3 into Level 4 maturity?

Level 4 requires proof that the customer has operationalized remediation. Better discovery is not enough. Better prioritization is not enough. The customer has to show that the organization can reduce risk through repeatable workflows.

The clearest proof is a before-and-after pattern: priority risks are identified, assigned, remediated, validated, and reported as reduced residual risk over time.

Outcome: Strong evidence includes closed-loop remediation records, SLA performance, exception governance, risk reduction metrics, cross-functional workflow adoption, DSPM/DLP coordination, audit-ready evidence, executive reporting, and repeatability across multiple repositories and teams.

Best-fit customer profiles

Best fit for organizations that have outgrown basic DLP and need a stronger data, AI and identity intelligence layer before they can operationalize risk reduction.

Best-fit customers include:

- Cloud and SaaS-forward enterprises with sensitive data spread across multiple repositories.
- Organizations with noisy DLP programs that need better classification, context, and policy tuning.
- Security teams trying to prioritize sensitive-data exposure by identity, access, sharing, and business context.
- Privacy, governance, and compliance teams that need better evidence of where regulated data exists and how it is controlled.
- Organizations preparing for AI adoption where sensitive data may be used by applications, models, agents, or unmanaged AI tools.

Poorer-fit customers include:

- Organizations looking only for a narrow endpoint DLP replacement.
- Organizations without the operating capacity to assign owners, act on findings, or manage remediation workflows.
- Buyers expect autonomous remediation (even if it's only a checkbox they won't use right away) before they have defined policy boundaries, approval models, and rollback requirements.

CISO implication: Customer profile criteria guide CISOs on whether their organizational maturity and infrastructure readiness justify a DSPM-led platform investment. If operational capacity is limited, CISOs should first focus on defining data ownership and approval workflows before procuring advanced automation platforms.

Competitive and adjacent category implications

Cyera's differentiation should be evaluated against adjacent categories, not only against other DSPM vendors.

The strategic implication is that Cyera's market power depends on whether buyers and adoptees see data security, identity and posture as a critical control plane for data security program and enforcement (DLP) modernization, AI and data protection, with identity-aware data access controls. If they see these as more loosely coupled areas, which SACR believes would be a critical mistake, they risk stalling out at earlier program maturity levels.

Comparing Cyera to:

- **Legacy DLP:** Cyera's advantage is context and a unified data model of data location and risk. The risk is that enforcement still depends on integration depth and existing control points, which is the compelling reason for DLP and DSPM convergence.
- **CNAPP:** Cyera's advantage is data sensitivity and business context. The risk is that some CNAPP platforms (and CSPs) may own cloud remediation workflows and buyer budgets.
- **Identity Security:** Cyera's advantage is data-aware access context. The risk is that identity platforms may own governance, access review, and entitlement remediation, [resolved by its acquisition of Oasis](#).
- **Data Governance Platforms:** Cyera's advantage is security posture and risk reduction. The risk is that governance buyers may prioritize catalog, lineage, privacy, and stewardship workflows over security remediation and automation as long-term goals.
- **AI Security Platforms:** Cyera's advantage is sensitive-data context and breadth of platform. The risk is that AI security markets continue to consolidate around runtime protection, model governance, or application security control points, and it doesn't keep up with this motion, mitigated by Cyera's browser, endpoint, gateway, and other inline control points.

Future Outlook for Data Security and Cyera

Cyera is useful as an example because it shows how a DSPM-led platform can attempt to connect data context, DLP modernization, AI governance, identity context, and remediation workflows. The buyer's question is not whether Cyera fits a category label, but whether it can prove customer movement from visibility to measurable risk reduction.

The next phase of data security will be less about discovering more data and more about governing the decisions that happen around sensitive data. That shift favors platforms that can connect data context to identity, exposure, movement, policy, remediation, and evidence.

Cyera's opportunity is to become a control layer for that transition. Fully autonomous Agentic remediation and policy changes are becoming an important differentiator, but only after the foundations are in place and only after customer trust is earned.

Exploitation of vulnerabilities is now the most common initial access vector for breaches, rising to 31% in the [2026 DBIR dataset](#), increasing the need to get security basics right, especially vulnerability management and remediation.

The near-term buyer value is better context, better prioritization, and better DLP decisioning. The longer-term value is governed automation that can safely reduce data risk at scale. To do so, the company must prove four things:

1. It can deliver high-confidence discovery and classification across the environments customers actually use.
2. It can improve DLP, access, and remediation decisions without forcing buyers into a rip-and-replace architecture.
3. It can help customers demonstrate measurable agentic-enabled risk reduction through operational workflows and autonomous or automated remediations, not just dashboards and graphs.
4. It can integrate and unify a single identity, data and control layer that can intercede, govern, control and prevent data leakage or misuse.

For security leaders, data-risk ownership is likely to keep moving out of security alone and into a shared model with data/identity/governance teams, regardless of which vendor they choose. Choosing a platform capable of uniting diverse stakeholders offers a stronger foundation for shared definitions of data sensitivity and policies capable of enforcing it.

Evidence buyers should request from Cyera:

CISOs and Practitioners should include the following during their evaluations:

- Which connectors and controls are generally available.
- How DSPM context improves DLP precision.
- Proof of reduced false positives / false negatives.
- Before-and-after remediation metrics.
- NHI / AI-agent governance use cases.
- Audit logs, approval gates, rollback, and validation evidence.
- Customer references showing movement from visibility to measurable risk reduction.

CISO implication: The future outlook highlights the need for CISOs to establish shared data-risk ownership models with privacy, identity, and business governance teams. CISOs should invest in trust-layer frameworks that supervise AI agent permissions, continuously auditing whether enterprise data security controls scale safely alongside AI adoption.

SACR takeaway

Cyera should be treated as a strong candidate in the data security platform conversation, particularly for customers moving through the layers of maturity and risk-driven DSPM and DLP alignment. The company's strategic upside is tied to the broader DLP reset and agentic control and data loss and access prevention, where legacy enforcement controls need richer data context, and DSPM platforms need to prove they can drive remediation rather than only discovery.

Methodologies

This research was sponsored by Cyera. SACR retained editorial control over the analysis and conclusions. Cyera briefings and factual review informed the report; roadmap, beta, and acquisition-integration claims were evaluated separately from generally available capabilities. SACR also evaluates data-security maturity through operating outcomes, evidence of generally available capability, customer-validated remediation impact, and alignment with CISO decision needs. Vendor briefings and factual review inform the report, but SACR retains editorial control. Roadmap claims, beta capabilities, and acquisition-integration plans are treated separately from generally available product capability.

Practical recommendations for CISOs

1. First 30 days

- Inventory sensitive-data coverage across SaaS, cloud, endpoints, collaboration, developer platforms, and AI tools.
- Identify top repositories with excessive access.
- Baseline DLP noise and unresolved high-risk findings.

2. Next 60 days

- Connect data context with identity and access context.
- Define ownership and SLA workflows for remediation.
- Build exception review and expiration discipline.

3. Next 90 days

- Pilot governed automation for low-risk, reversible actions.
- Validate remediation outcomes.
- Report maturity using outcome metrics, not tool deployment status.

Appendix: CISO Data Maturity Model Evaluation Checklist

Maturity Self-Assessment: Evaluating Your Data Security Journey

Use this practical checklist to evaluate your organization's current data security maturity across four critical dimensions. Match your operational reality to the descriptions below to determine whether you are operating at Level 1 (Initial), Level 2 (Repeatable), Level 3 (Defined), Level 4 (Managed), or Level 5 (Optimized).

Suggested Questions every CISO should ask:

1. Where does our sensitive data live across SaaS, cloud, endpoints, collaboration tools, developer environments, and AI workflows?
2. Who or what can access that data, including users, service accounts, APIs, automation, and AI agents?
3. Which exposures create the greatest business, regulatory, or operational risk?
4. Which controls can act on this context today? Can they warn, block, quarantine, encrypt, escalate, remediate, or validate?

5. How do we prove that remediation reduced risk rather than merely closed a ticket?
6. What automation is allowed, what requires approval, and what must remain human-led?
7. Which Cyera capabilities are generally available today, and which depend on roadmap integration or services?

CISO Checklist Items

1. Discovery & Data Inventory

- **Level 1 (Initial):** Are your data inventories maintained manually via static spreadsheets, leaving significant blind spots across cloud and SaaS environments?
- **Level 2 (Repeatable):** Do you run scheduled, batch classification scans across core data stores, though visibility remains delayed and fragmented across newer data stores?
- **Level 3 (Defined):** Do you have automated data discovery that covers major cloud repositories, SaaS apps, and databases with systematic data labeling?
- **Level 4 (Managed):** Is data discovery integrated with identity, access, and entitlement context to continuously highlight toxic combinations of exposure and sensitive data?
- **Level 5 (Optimized):** Does discovery occur at ingestion time (e.g., upload/commit), preventing unvetted data from entering the environment while maintaining continuous lifecycle lineage?

2. Policy & Enforcement

- **Level 1 (Initial):** Are policies reliant on basic regex or legacy pattern matching, generating high false-positive rates (>50%) that exhaust security analyst bandwidth?
- **Level 2 (Repeatable):** Are DLP policies documented and applied consistently across primary communication channels (email, endpoint, network), though tuning remains heavily manual?
- **Level 3 (Defined):** Do policies incorporate basic risk scoring and API-level access controls to enforce boundary-level controls dynamically?

- **Level 4 (Managed):** Are policy engines context-aware, incorporating user behavior, non-human identity (NHI) activity, device posture, and threat intelligence into real-time decisions?
- **Level 5 (Optimized):** Does your enforcement architecture feature automated remediation, adaptive policy adjustment with predictive threat detection and minimal false positives (<1%)?

3. Incident Triage & Remediation Workflow

- **Level 1 (Initial):** Is incident response purely reactive, driven by manual ad-hoc investigations after an event or breach notice?
- **Level 2 (Repeatable):** Are alerts logged centrally and routed via standard ticketing systems, though >80% of remediation work remains manual and slow?
- **Level 3 (Defined):** Are clear data owners assigned to findings with standardized remediation guidance and SLA tracking for high-priority exposure risks?
- **Level 4 (Managed):** Do you utilize closed-loop remediation workflows with automated exception handling, SLA tracking, and post-remediation verification?
- **Level 5 (Optimized):** Are agentic workflows capable of autonomous, governed remediation—such as instantly revoking risky exposure or killing compromised service account sessions at machine speed?

4. Governance & Executive Measurement

- **Level 1 (Initial):** Is compliance reporting performed ad hoc to satisfy immediate audit requests, with no ongoing posture measurement?
- **Level 2 (Repeatable):** Do you produce periodic compliance reports and basic risk dashboards for executive leadership based on point-in-time assessments?
- **Level 3 (Defined):** Are security controls systematically mapped to regulatory frameworks (e.g., GDPR, HIPAA) with defined risk-tolerance thresholds?
- **Level 4 (Managed):** Can you provide board-level reporting on residual data risk reduction, SLA adherence, and cross-domain posture metrics over time?
- **Level 5 (Optimized):** Is continuous, audit-ready evidence automatically generated alongside continuous optimization metrics to demonstrate dynamic alignment with business risk tolerance?



business



personal



softwareanalyst.io



Trusted research. Sharp insights. Real conversation.

CISO

VENDOR

SECURITY
TEAMS

INVESTORS