COHESITY

BUYER'S GUIDE

NAVIGATE DATA AND AI RISK WITH THE RIGHT DSPM SOLUTION

Executive summary

Sensitive data is now distributed across your cloud platforms, SaaS, databases, backups, hybrid infrastructure, NAS, and more. You likely already have more data than you can easily govern. But it's hard to understand where your most sensitive data lives, how broadly it's exposed, who has access to it, and whether that access is appropriate.

This creates a perfect storm that amplifies business risk at the board level. And the challenge is growing: Even as people continue accessing, altering, and moving this data, AI agents are introducing new and growing risk.

These security and protection gaps are why Data Security Posture Management (DSPM) has become a strategic priority for modern enterprise leaders.

89% **of security leaders indicate that the data security status quo at their company is a problem.**

39% **of security leaders indicate that legacy technologies are insufficient for current requirements.**

60% **of enterprises lack visibility into at least half of their data estate.**[1]

[1]Omdia Whitepaper: Achieving Cyber Resilience With Actionable Data Intelligence

A modern DSPM solution should help you do far more than inventory data. It should help you actively reduce risk, strengthen cyber resilience, and answer critical questions during incident response, security drills, or audits:

- Where and what is our sensitive data?
- Who or what has access to it?
- What should we fix first?
- How do we reduce risk without slowing business innovation and growth?
- Is the data protected and can it be recovered quickly, securely, and compliantly?

For security leaders, this means moving from visibility to actionable risk reduction both before and after an attack occurs.

This guide is for teams evaluating DSPM solutions, so you can select the right one for your needs.

Inventory is not enough.

The real buying question isn't whether a solution can show you where the data is. It's whether the solution can identify and show *why* it's high risk and what you can do now to remediate that risk.

What DSPM solves

One of the most common mistakes organizations make when evaluating DSPM is treating it like a simple discovery tool. A DSPM solution is about much more than visibility. The right solution should help you solve problems around data risk posture.

Here are five problems your DSPM solution should help you solve:

Problem	How DSPM solves it
I don't know where my sensitive data is.	Many organizations cannot confidently answer where their highest-value or regulated data resides across cloud, SaaS, and hybrid environments. Why is this the status quo? Traditional data security tools are siloed to specific types of data, certain deployment models, or specific use cases. Even when they have some inventory, it is often incomplete, outdated, or disconnected from risk context. A modern, cloud-native architecture is required to keep pace with the volume and sprawl of data.
I don't know who or what is accessing (or can access) my data.	The typical employee interacts with roughly 36 cloud services each day, across more than 200 collaboration tools. Sensitive information spreads across platforms and is often shared outside established governance controls. Sensitive data becomes even riskier when it is overexposed, over-permissioned, shared too broadly, stored in the wrong place, or accessible to more identities (and non-human ones) than intended. Your DSPM solution should help you understand which data is at risk, not just which data matches a classification policy.
I don't know what data risks I should prioritize first for remediation.	Most enterprise environments contain large amounts of sensitive data. However, not all data risks carry equal urgency. Teams need a way to distinguish background noise from meaningful exposure. The right solution should help you focus on the issues with the greatest potential business impact, based on sensitivity, access, exposure, environment, and learned business context.
I can't remediate data risks quickly or automatically.	Data context—combined with deep integration across the security toolchain—is essential for reducing MTTR (mean time to respond) for remediating data that is at high risk. Traditional data security methods rely heavily on patterns and tags, flooding teams with noisy alerts that often lack credibility. DSPM solutions should drive action. Findings should be understandable, assignable, and trackable. Security and IT teams should be able to act on them with confidence.
When I need to recover data, it's slow, risky, and costly.	Secondary data (such as backups) is critical when it comes to incident response and recovery. This data can also introduce risk and potential compliance violations that need to be addressed. Modern DSPM solutions should strengthen cyber resilience outcomes and feed data intelligence to help ensure the right levels of protection are applied. Sensitive data should be prioritized and recovered compliantly during incident response.

What "good" looks like

Before evaluating vendors, you need to decide what "good" looks like. Here are five areas where a modern DSPM solution should excel:

Deploy quickly and scale seamlessly without limits

Good discovery should help you quickly gain a more complete view of your exposure without requiring a long, heavy implementation cycle before value appears.

Organizations should expect continuous scanning of hundreds of petabytes, without compromising performance. The right vendor will offer broad support across key environments, including cloud data stores, SaaS apps, file systems, and on-prem databases where applicable.

Eliminate blind spots with accurate discovery and classification

Strong discovery helps prevent inferior results from contaminating every downstream outcome. If the DSPM tool misses critical locations, misclassifies sensitive information, or requires too much manual tuning, the findings will be less helpful, or even irrelevant.

The right DSPM solution uncovers blind spots by adapting to your environment and classifying data automatically (including IP and data that is uniquely important to your business).

Prioritize real risk, not noise

Time is limited; no team can remediate everything at once. A strong DSPM solution combines exposures and context to focus on the biggest dangers. It should prioritize risk from highly sensitive data with broad exposures.

This requires advanced capabilities to correlate data sensitivity, business purpose, identities, access activity, and exposure to surface real business risk.

Automate remediation at scale

Modern DSPMs need to help organizations reduce exposure, not just observe it. Organizations should look for DSPMs that provide out-of-box remediations. For example, leading solutions offer one-click controls to revoke access, mask data, trigger predefined workflows, and route issues directly to data owners, all with the context needed to resolve issues quickly.

This is particularly important at enterprise scale, when many teams are involved. Security teams often identify the issue, but remediation may depend on different groups like storage administrators, data owners, identity teams, or business stakeholders. The more clearly a solution supports that handoff, the more likely it is to drive meaningful outcomes that improve your security posture.

Strengthen security before and after incidents

We've already covered proactive risk reduction before a breach occurs. The right DSPM solution can also provide the data intelligence and context needed to strengthen recovery efforts, ensuring security and data protection teams understand where sensitive data resides, which assets are most critical, and whether the appropriate protections have been applied to the organization's crown jewels.

With clear insight into sensitive and business-critical data, your DSPM should inform and enhance the data protection strategies that underpin cyber resilience. This will help you prioritize the most important data during recovery, accelerate incident response with better context, and reduce cost and complexity by identifying and eliminating stale, redundant, or low-value data.

The result is a continuous feedback loop between visibility, protection, response, and recovery. With the right DSPM solution, you should be able to connect these functions so you can move from siloed security efforts to a more holistic, end-to-end data security and cyber resilience strategy.

Vendor evaluation kit

Critical questions to ask vendors

The following questions will help you evaluate your shortlist of DSPM products.

On discovery and coverage:

- Which cloud, SaaS, AI, database, and hybrid environments do you support today?
- How do you identify and classify sensitive data across different data types? How long does it typically take?
- How do you handle structured and unstructured data? Do you have better support for one or the other?
- What level of tuning is required before results are usable?
- How do you help customers validate coverage and confidence early?

On exposure and access controls:

- How do you determine whether sensitive data is overexposed? Can you show which identities, groups, or users have access to sensitive data?
- How do you identify broad, excessive, or risky access?
- How do you connect access context to data sensitivity?
- Can you block sensitive data from being entered or leaving AI models and systems?
- Can you show how your solution highlights the issues of greatest risk?

On risk prioritization and remediation:

- How do you rank or prioritize findings? What factors influence severity or risk scoring?
- Can you show how teams move from many findings to a manageable action plan?
- How does your solution guide remediation?
- Can findings be assigned to the right teams or owners?
- Can insights be fed to data protection platforms seamlessly and provide better security and recovery outcomes?

On time to value and deployment:

- What does initial deployment require from our team?
- How quickly can we expect useful findings?
- What should success look like in the first 30, 60, and 90 days?

Red flags to look out for

Limited visibility into access risk

Access to sensitive data is the biggest risk. If a vendor can't clearly show who (or what) can access data and what risks exist—this is a serious limitation.

Findings aren't connected to individual business risk and context

Every business is different. The risk posed by sensitive data or exposed intellectual property can vary based on what's important to that organization. Findings should be based on learned context and classifiers to deliver accurate risk remediation recommendations.

Inability to tie data risk to cyber resilience outcomes

Optimizing data and AI risk posture is critical for cyber resilience. Data intelligence should be used to inform proactive data security and enable better recovery outcomes. Findings should not be siloed, especially since remediation often involves cross-functional teams.

Value depends on excessive manual effort

A product that requires extensive tuning, interpretation, or customer process design can add ongoing operational cost and complexity.

Vendor comparison table

Use this table to compare vendors side by side across the core capabilities required for an effective DSPM strategy. As you review each row, consider not just whether a feature exists, but how well it is delivered in practice*.

To make vendor comparisons more consistent and objective, each capability can be evaluated using a simple 0 to 5 scoring model. This approach helps teams move beyond a basic check-the-box comparison and assess the relative strength, maturity, and practical usability of each capability. A higher score should indicate that the capability is not only available, but also enterprise-ready, operationally efficient, and proven to deliver value at scale.

Score	Definition
5	Best-in-class capability that is fully delivered, highly automated, scalable, and well integrated
4	Strong capability with solid functionality and minor gaps
3	Adequate capability that meets core requirements but has limitations in depth, scale, or usability
2	Basic or partial support with notable gaps, manual effort, or limited applicability
1	Minimal capability, early-stage functionality, or highly constrained support
0	Not supported or no meaningful evidence provided

Category	Vendor 1	Vendor 2	Vendor 3	Notes
Cloud-native architecture for quick and agentless deployment				
Continuous data discovery at scale and speed				
AI-powered data classification				
Data contextualization				
Risks mapped to regulations				
Adaptable access controls and enforcement				
Real-time data detection and response				
Contextualized alerts				
Integrated remediation workflows				
Seamless integration with data protection platforms				
Total				

**Scores should be based on what is generally available today, not roadmap claims. Each score should be validated using product demos, customer references, documentation, deployment requirements, and proof of integration depth.*

The bottom line

Selecting a DSPM solution is more than a simple tooling decision. The stakes are too high. Data continues to proliferate and sprawl at an alarming rate, and manual approaches or legacy technology cannot keep pace. The right DSPM solution will help you create clarity, focus your teams, reduce risk, and strengthen both data security and cyber resilience outcomes. This solution will pay dividends in the long run, as it will help you protect and secure the data that underpins AI models, systems, and agents.

Learn more about our DSPM solution and how it can optimize your data and AI risk posture. Or contact us today.

COHESITY

cohesity.com
1-855-926-4374
2625 Augustine Drive, Santa Clara, CA 95054

4500066-001-EN 4-2026