Which DSPM vendors are actually worth evaluating today?

A current, evidence-led buyer's guide for enterprise security, data-security, cloud-security, privacy, and governance leaders — researched 29 July 2026.

Executive answer

There is no defensible universal "best DSPM." There is, however, a defensible shortlist.

For a serious enterprise evaluation today, start with **Cyera, BigID, Varonis, Sentra, and Wiz**. Those five expose the most useful differences in the market: a modern data-security pure play; broad enterprise discovery and governance; deep permissions and activity-centric security across hybrid collaboration estates; a cloud-first specialist; and DSPM embedded in a cloud-native application protection platform (CNAPP). Add **Microsoft Purview** when Microsoft 365, Copilot, Fabric, and the existing Purview control plane dominate the requirement. Add **Securiti, now part of Veeam**, when privacy, data governance, AI governance, and resilience must converge. Add **Proofpoint DSPM (formerly Normalyze)** when human-centric risk, DLP, lineage, and attack-path analysis are central.

Four more vendors are worth a conditional look:

- **Palo Alto Networks Cortex Cloud DSPM** for organizations already standardizing cloud security operations on Palo Alto Networks.
- **Rubrik DSPM (formerly Laminar)** when cloud-data posture must connect to backup, cyber recovery, and Microsoft 365/Copilot readiness.
- **Thales CipherTrust DSPM** when discovery must lead to native encryption, tokenization, masking, and key management.
- **Concentric AI** when unstructured data, semantic classification, least-privilege cleanup, and data in motion are the hard part.

IBM Guardium also belongs in a regulated-enterprise or incumbent-Guardium request for proposal (RFP), although the public product evidence available for its current DSPM packaging is less clear than for the leading shortlist. **Tenable, OpenText, Skyhigh Security, and other suite vendors can be credible in their installed bases**, but should not displace a specialist benchmark in a greenfield selection.

That conclusion is deliberately not a league table. The category now spans five different use cases—DLP, privacy and governance, entitlement management, cloud-posture integration, and AI-workflow protection—and products do not cover them equally. A May 2026 *CSO* review, drawing on Gartner and other analyst research, makes exactly that point and warns that connector coverage, integration depth, and separately licensed controls require close examination (CSO, "DSPM buyer's guide," 27 May 2026). Omdia's 2025 comparative assessment likewise found leaders and challengers with materially different strengths, and explicitly noted that most offerings still had gaps, particularly in reactive remediation (Omdia Universe: DSPM, 2025).

The practical recommendation is therefore:

- Pick **three vendors from different architectural families**, not three near-identical products.
- Make them scan the same representative estate with the same permissions.
- Score verified coverage, classification quality, effective-access accuracy, operational cost, and closed remediation—not slideware.
- Treat every classifier count, "petabyte scale," accuracy percentage, time-to-value assertion, and ROI claim as a hypothesis until reproduced on your data.

What "worth evaluating" means

A vendor is worth evaluation if public evidence shows a coherent product, a plausible operating model, relevant source coverage, and a route from discovery to reduced exposure. It does **not** mean the vendor has been independently proven to be most accurate, cheapest, or safest. Public DSPM evidence is unusually asymmetric: product documentation and customer stories are plentiful; reproducible cross-vendor benchmarks, blinded classification studies, false-negative measurements, and transparent price books are scarce.

This article uses four evidence classes:

- **Independent comparative evidence**, principally Omdia's 2025 assessment and *CSO*'s May 2026 buyer's guide.
- **Primary product documentation**, used to establish what a vendor says its product supports and how it is designed.
- **Primary corporate records**, used to verify acquisitions and current ownership.
- **Vendor customer stories**, used only as documented examples—not as neutral proof of general performance.

Vendor statements are described as claims. Analyst rankings are treated as analyst judgments, not laboratory results. No public source reviewed provides a controlled, current, apples-to-apples benchmark of discovery recall, classification precision and recall, effective-access accuracy, scan overhead, or total cost across the shortlisted products.

DSPM is no longer one product category

At minimum, a DSPM should continuously answer four questions: **what sensitive data exists, where it is, who or what can reach it, and whether it is appropriately protected**. Mature products add usage and lineage, contextual risk, policy, guided or automatic remediation, evidence for compliance, data detection and response (DDR), and controls for AI systems.

The market's expansion explains why superficially similar products behave differently:

Architectural family	Typical center of gravity	Natural advantage	Typical risk to test
Data-security pure play	Data discovery, classification, access graph, risk	Cross-platform data context and rapid innovation	Protection may depend on integrations or extra modules
Data governance/privacy platform	Catalog, privacy, retention, lineage, policy	Broad sources, ownership, compliance workflows	Security operations and real-time response may be less mature
Permissions/activity security platform	Effective access, behavior, least privilege	Actionable exposure reduction and threat detection	Deepest support may cluster around selected SaaS/file ecosystems
CNAPP-embedded DSPM	Cloud graph, identity, workload, exposure paths	Correlation with vulnerabilities, identities, and attack paths	SaaS, on-premises, privacy, and lifecycle depth may lag
Data protection/resilience suite	Backup, recovery, data inventory	Connect posture to recoverability and cyber resilience	Production discovery and remediation breadth must be proved
Native ecosystem control plane	One hyperscaler or productivity suite	Low-friction policy enforcement in the incumbent estate	Cross-cloud breadth may be integration-led rather than native

This is not semantic hair-splitting. *CSO* notes that products may be SaaS, on-premises, private-VM/container, or hybrid; that "agentless" normally means API or metadata-oriented access; that scan cadence varies; and that every vendor's claim to support AWS, Azure, and Google Cloud does not imply support for every relevant service (CSO). Buyers should therefore compare **specific source × data type × deployment mode × security-context × remediation-action combinations**, not logo walls.

The core shortlist

1. Cyera: the benchmark modern pure play

Evaluate when: the enterprise wants a data-security control plane across cloud, SaaS, databases, AI, and on-premises sources, with classification and remediation as the primary buying center.

Cyera is the most useful pure-play benchmark because its product proposition is centered on DSPM rather than inherited from an adjacent category. Its current product page says the platform covers cloud, SaaS, DBaaS, and on-premises stores; uses agentless scanning; supports structured and unstructured information; and offers more than 30 out-of-the-box actions such as revoking access, masking data, and routing work to owners (Cyera DSPM). *CSO* reports agentless visibility across structured, semi-structured, and unstructured data, more than 500 built-in classifiers, security/data ecosystem integrations, and on-premises or hybrid deployment options for in-environment scanning and residency needs (CSO).

The product's appeal is breadth with a data-first graph: discover data, classify it, join it to identities and exposure conditions, then drive actions. Cyera's adjacent Access Trail, data-loss prevention, DDR, and AI-security products also show a credible path beyond a passive inventory. That makes Cyera particularly relevant to organizations seeking one specialist platform that can become the system of record for data risk.

The caution is evidence quality. Cyera advertises "95%+ precision," scanning at hundreds-of-petabytes scale, rapid deployment, and customer outcomes on its own page (Cyera); the public page does not provide a reproducible test set, confusion matrix, sampling design, or independent validation for those figures. Precision alone also says nothing about recall: a classifier can be precise while missing important sensitive data. Test both. Buyers should also identify which remediation actions are truly native, which call third-party APIs, which require another Cyera module, and what rollback and approval controls exist.

Proof-of-value tests

- Measure precision and recall on a labeled corpus containing organization-specific identifiers, multilingual text, tables, scans, source code, and deliberately ambiguous near-matches.
- Compare effective access—not merely direct permissions—for nested groups, public links, service principals, cross-account roles, inherited ACLs, and dormant identities.
- Confirm what content or samples leave the customer boundary, where metadata is processed, how it is retained, and how regional isolation works.

- Price initial scan, continuous rescanning, SaaS connectors, on-premises coverage, AI modules, DLP/DDR, and remediation separately.

Bottom line: Cyera belongs in most greenfield enterprise evaluations, but its marketing numbers must be treated as claims to validate, not selection evidence.

2. BigID: broad discovery, governance, privacy, and security in one data platform

Evaluate when: the organization has a highly heterogeneous estate, significant privacy/governance obligations, legacy and on-premises sources, or needs security to share a data-intelligence layer with privacy and data offices.

BigID's differentiator is breadth. Its current product material claims structured, semi-structured, and unstructured discovery across cloud, SaaS, IaaS, PaaS, development tools, applications, email, big-data/NoSQL, on-premises, hybrid, and air-gapped environments, with more than 500 native source connections (BigID DSPM and AI-SPM). Its main DSPM page describes discovery, ownership/access context, risk prioritization, and actions including deletion, redaction, access revocation, retention enforcement, and owner-delegated remediation (BigID DSPM).

There is independent support for taking it seriously. Omdia placed BigID among four leaders in its 2025 assessment and said it delivered the broadest overall capability portfolio in that comparison (Omdia). BigID's security proposition also sits beside privacy, catalog, lifecycle, and AI-governance functions. For an enterprise already trying to reconcile CISO, chief data officer, privacy, and records-management inventories, that convergence can be more valuable than a narrowly optimized cloud scanner.

The same breadth is the evaluation risk. A large platform can require more architecture, tuning, stewardship, and module selection than a focused product. BigID's public claims—such as "1,500+ classifiers," "95% faster than legacy tools," and unique agentic remediation—are vendor assertions without a public neutral benchmark on the cited pages (BigID). Buyers must also distinguish *connector existence* from deep support: can the connector discover all relevant objects, inspect content without breaking controls, resolve effective access, ingest activity, understand lineage, and remediate safely?

Proof-of-value tests

- Select ten difficult sources, not ten popular ones, and grade each across inventory, content inspection, permissions, activity, lineage, ownership, and write-back.
- Require an implementation plan showing scanners, network paths, credentials, data residency, high availability, upgrades, and operational ownership.
- Run the same custom-classifier creation task in BigID and a focused specialist; measure analyst effort and maintainability as well as accuracy.
- Demonstrate one privacy workflow, one security exposure workflow, one records/lifecycle workflow, and one AI-data workflow end to end.

Bottom line: BigID is one of the strongest candidates where coverage and organizational convergence matter more than minimal deployment footprint.

3. Varonis: strongest fit for permissions, usage, and active exposure reduction

Evaluate when: oversharing, stale access, identity risk, SaaS collaboration data, file shares, and anomalous data activity are the primary problems—especially across Microsoft 365, Google Workspace, Salesforce, and hybrid file estates.

Varonis predates the DSPM label and approaches the problem through data access, permissions, and behavior. Its current material says the platform monitors SaaS, IaaS, and file storage; manages third-party app risk; records forensic data-access and permissions activity; baselines behavior; and can intervene in risky access (Varonis SSPM and data security). *CSO* highlights cloud and on-premises coverage, integrations across SIEM, SOAR, network, identity, and Microsoft information-protection systems, and a managed data detection and response service using behavioral models and automatic remediation (CSO).

Omdia categorized Varonis as a challenger in 2025, while noting strong strategy/innovation, advanced features, and solution breadth (Omdia). The label matters less than fit: a tool that can explain effective access, actual usage, and safely remove excessive access may create more security value than one that finds a larger raw inventory but hands remediation to another team.

The evaluation question is coverage balance. Varonis's heritage and public evidence are especially strong around unstructured collaboration/file data and permissions. Buyers with a database-, lakehouse-, engineering-, or multicloud-PaaS-heavy estate should validate every required platform and data type against a data-discovery specialist. Also separate product-native automation from the managed service, and determine what telemetry is continuous versus periodic.

Proof-of-value tests

- Seed excessive access through nested groups, organization-wide links, external guests, stale accounts, service identities, and OAuth applications; ask Varonis to calculate the true blast radius.
- Require activity-led prioritization: identify sensitive assets that are both broadly reachable and actively accessed, and distinguish them from dormant theoretical exposure.
- Automate a least-privilege change in a test tenant, with owner approval, exception handling, rollback, and a complete audit trail.
- Test database, object-store, code, and AI-workload coverage as rigorously as Microsoft 365 and file shares.

Bottom line: Varonis is a must-evaluate where "who can and does touch sensitive data?" is more important than "how many stores can we catalog?"

4. Sentra: focused cloud-first DSPM with an increasingly broad security workflow

Evaluate when: the estate is cloud-first or data-platform-heavy and the buyer wants a specialist that scans in place, emphasizes classification context, and adds near-real-time detection without buying a full CNAPP.

Sentra says it scans data in place and keeps customer data inside the customer cloud (Sentra DSPM definition). *CSO* reports support across cloud services, containers, virtual machines, on-premises sources and file shares, plus DDR, integrations across ITSM, identity, SIEM, CNAPP, collaboration, and data-management tools, and automated remediation (CSO).

Omdia described Sentra's 2025 proposition as balanced and innovative but noted that its smaller size limited overall breadth and scope relative to larger competitors (Omdia). That is a fair framing: Sentra is useful in a bake-off precisely because it tests whether a focused, modern cloud-data architecture outperforms a sprawling suite on the sources that matter.

The risks to evaluate are depth outside the cloud center of gravity, operational maturity across every connector, and the boundary between posture scans and "near-real-time" DDR. "In-place" also needs precise architectural confirmation: where does classification execute, what metadata or samples are returned to the SaaS control plane, and what temporary copies or snapshots are created?

Proof-of-value tests

- Benchmark production-scale Snowflake/Databricks/cloud database and object-store scans for compute cost, query impact, snapshot cost, throttling, and freshness.
- Generate an access or activity event and measure detection latency and context completeness.
- Test how classifiers propagate across copies and whether lineage survives transformations.
- Compare coverage of the buyer's three hardest SaaS or on-premises sources with BigID or Varonis.

Bottom line: Sentra deserves a place in cloud-first evaluations and is a useful specialist counterweight to CNAPP bundles.

5. Wiz: best benchmark for DSPM inside a CNAPP/cloud graph

Evaluate when: cloud security owns the problem; the organization already uses or is considering Wiz; and the priority is connecting sensitive data to public exposure, identities, vulnerabilities, malware, workloads, code, and attack paths.

Wiz's product page describes continuous agentless discovery across code, IaaS, PaaS, DBaaS, and AI, with built-in and custom classifiers; it correlates sensitive data with public exposure, identities, entitlements, vulnerabilities, malware, and lateral movement in the Wiz Security Graph (Wiz DSPM). Wiz has also documented extending its classifiers into code bases, pull requests, IDEs, and CI/CD platforms (Wiz, "extends DSPM capabilities to code bases," 12 March 2024).

This is the strongest reason to evaluate Wiz: the product can place "database contains sensitive data" inside a plausible cloud attack path, instead of presenting it as an isolated data finding. For cloud-platform and application-security teams, that context can improve prioritization and let one workflow reach developers and infrastructure owners.

The trade-off is the CNAPP boundary. Wiz's public DSPM proposition is cloud-centric. Buyers whose critical risk sits in Microsoft 365, Google Workspace, Salesforce, enterprise file shares, privacy workflows, retention, or deep database activity should compare it with a data-security specialist. Also establish which remediation is native, which is ticket/orchestration-driven, and which requires other Wiz capabilities.

Wiz is now part of Google: *CSO*'s May 2026 market summary records Google's acquisition of Wiz (CSO). That may increase platform resources, but buyers should ask for a contractual roadmap for AWS and Azure neutrality, packaging, support, and telemetry boundaries rather than assuming either benefit or harm.

Proof-of-value tests

- Construct a test attack path involving an internet-reachable workload, exploitable weakness, overprivileged identity, and sensitive data; compare Wiz's graph with a pure-play DSPM plus the existing CNAPP.
- Test code and CI/CD findings for secrets and sensitive sample data, including developer routing and suppression lifecycle.
- Validate non-Google-cloud parity and obtain roadmap commitments for critical AWS/Azure services.
- Compare the additional cost and operational benefit of the DSPM module with importing specialist DSPM findings into Wiz.

Bottom line: Wiz is a compelling choice for cloud-centric attack-path prioritization, but it should not automatically be treated as an enterprise-wide data-security platform.

Add these vendors when the operating model fits

Microsoft Purview: the incumbent-first option for Microsoft data and AI

Microsoft's current Purview DSPM documentation describes unified visibility and controls spanning Microsoft 365, Azure, Fabric, AI apps/agents, and integrated third-party SaaS platforms. It consolidates signals from Purview DLP, Insider Risk Management, Information Protection, and Data Security Investigations; supports objectives and recommended actions; and can remove public links or apply policies. Microsoft also explicitly lists integrations with Varonis, Cyera, BigID, and OneTrust for broader risk insight (Microsoft Learn, "Learn about Data Security Posture Management," updated 2026).

That makes Purview a serious evaluation—not merely a licensing checkbox—when Microsoft 365, Copilot, Entra, endpoints, and Purview labels already define the control plane. Native policy enforcement and existing classification investments can outweigh a specialist's broader inventory.

The limitation is equally clear in Microsoft's documentation: Microsoft locations in the Asset Explorer currently include Microsoft 365, while some non-Microsoft locations depend on partner integrations; setup for some non-Microsoft data sources is described as preview (Microsoft Learn). Microsoft also warns that generative-AI investigation results may be inaccurate or incomplete and should be verified (Microsoft Learn, Data Security Posture agent). Buyers must distinguish the new DSPM from "classic" versions and confirm preview status, licensing, regional availability, data prerequisites, and partner dependency.

Verdict: shortlist Purview for a Microsoft-centered estate; pair it against a specialist for heterogeneous coverage.

Securiti, now Veeam: convergence of DSPM, privacy, governance, AI trust, and resilience

Securiti was already a credible broad platform before the transaction. Omdia placed it among challengers in 2025 and rated its advanced capabilities best in class in that comparison (Omdia). Veeam announced the completed $1.725 billion acquisition on 11 December 2025 and described the combination as spanning DSPM, privacy, governance, AI trust, production data, secondary data, and resilience (Veeam acquisition completion).

This creates a distinctive evaluation hypothesis: one inventory and policy plane could govern data and AI while also understanding backup and recovery. That is attractive to organizations trying to unify CISO, privacy, governance, and resilience programs.

Acquisition integration is also the principal risk. Require a current architecture—not a future-state diagram—showing identity, consoles, data models, policy, deployment, support, and licensing. Ask which workflows are genuinely unified today and obtain dated commitments for the rest.

Verdict: evaluate for broad data/AI governance plus resilience, but score shipping integration separately from roadmap.

Proofpoint DSPM, formerly Normalyze: human risk, lineage, and data attack paths

Proofpoint's current page says its Normalyze-derived DSPM classifies structured and unstructured information in cloud and on-premises environments, prioritizes unauthorized or over-permissioned access, and supports recommended or automated workflows (Proofpoint, "Normalyze is now Proofpoint"). Normalyze also documented a one-pass scanner, data-access graphs, an attack-path-oriented Risk Navigator, relative data valuation, custom risk signatures, and data-in-motion/lineage capabilities (Proofpoint DSPM guide; Normalyze lineage announcement).

The strategic case is the join between data posture and Proofpoint's human-centric security, DLP, email, insider-threat, and cloud controls. *CSO* reports more than 300 classifiers, integrations with ITSM/orchestration tools and Purview, and protection for AI workflows (CSO).

As with Veeam/Securiti, the acquisition makes integration maturity the key diligence item. Validate whether shared identity, policy, incidents, classifiers, telemetry, agents, administration, and licensing are unified or simply linked.

Verdict: a strong conditional candidate where users, data movement, DLP, and attack paths must meet.

Conditional specialists and suite choices

Palo Alto Networks Cortex Cloud DSPM

Palo Alto Networks acquired Dig Security and embedded the technology in its cloud platform. Current field guidance describes an agentless DSPM module and explicitly notes that the Dig-derived module is deployed to a tenant (Prisma Cloud Field Guide). *CSO* reports hundreds of SIEM/workflow/ticketing integrations, more than 600 prebuilt classifiers, and coverage including Microsoft 365, Snowflake, SaaS, clouds, and on-premises file shares (CSO).

Worth it when: Palo Alto is the strategic SOC/cloud platform and reducing tool fragmentation has real operational value.

Prove: parity between inherited Dig capability and current Cortex packaging; cross-cloud/source depth; classifier quality; and whether remediation, XDR correlation, and AI controls work in one incident workflow.

Rubrik DSPM

Rubrik's DSPM originates in Laminar. Omdia said Rubrik had successfully integrated it with backup and recovery, giving a combination of proactive posture and reactive breach remediation, although additions to core technology would improve its score (Omdia). Rubrik's own guidance says mature DSPM should discover known and unknown data, support hygiene, access governance, compliance, and contextual remediation, scan inside the customer environment, minimize redundant rescans, and handle structured, semi-structured, and unstructured data (Rubrik DSPM guide). Rubrik also documents Microsoft 365 Copilot posture functions including discovery, risky-permission removal, and Purview label correction (Rubrik DSPM for Microsoft 365 Copilot).

Worth it when: sensitive-data discovery, backup coverage, cyber recovery, and Microsoft 365 readiness should converge.

Prove: production-store breadth independent of protected backup assets, identity/access depth, actionability, and the integration of the Laminar plane with Rubrik Security Cloud.

Thales CipherTrust DSPM

Thales is unusual because it can connect discovery to native cryptographic controls. Omdia named Thales a 2025 leader, citing core technology, momentum, execution, discovery, encryption, tokenization, masking, keys, secrets, identity, and centralized control. Omdia also identified limitations: less comprehensive classification for some file types and no DSPM-native incident-response planning, with some functions on the roadmap or supplied through the broader portfolio (Omdia).

Worth it when: regulated data protection, sovereign key control, encryption, tokenization, HSMs, and masking are procurement-critical.

Prove: the operational join between Imperva-derived discovery and CipherTrust protection; classifier coverage for the actual formats; and whether remediation is a coherent workflow or a set of adjacent products.

Concentric AI

Concentric positions its Semantic Intelligence platform around context-aware AI that discovers, categorizes, monitors, and protects data at rest and in motion across cloud and on-premises environments (Concentric AI). Omdia rated the smaller vendor as a challenger with top-tier solution breadth and a strong advanced posture-management proposition (Omdia).

Worth it when: unstructured documents, semantic business context, data access governance, and least-privilege remediation dominate.

Prove: accuracy on organization-specific unstructured material; explainability; multilingual and image/OCR performance; SaaS/database depth; and data-in-motion enforcement latency.

IBM Guardium

Omdia placed IBM among its four 2025 leaders, noting that IBM acquired Polar Security in May 2023 and integrated DSPM into Guardium (Omdia). This is enough to warrant evaluation for existing Guardium customers and highly regulated enterprises that value database-security monitoring and a large-vendor operating model.

The reservation is not that Guardium lacks capability; it is that current, granular, publicly accessible DSPM documentation was harder to verify than for the primary shortlist. Require a live source matrix, architecture, deployment and licensing guide, release history, and customer references matching the intended estate.

Verdict: include in incumbent and regulated-enterprise RFPs; do not waive a specialist bake-off.

Vendors not on the default shortlist

Exclusion is not a declaration that a product is bad. It means the public evidence does not justify putting it ahead of the core candidates for a general greenfield enterprise evaluation.

- **OpenText** was an Omdia 2025 leader with best-in-class strategy and solution breadth (Omdia). It is reasonable for an OpenText-installed base, but a buyer should demand a crisp demonstration of the integrated current cloud platform and compare usability and deployment effort with a specialist.
- **Skyhigh Security** earned strong advanced-capability scores in Omdia's assessment but was categorized as a market prospect, with Omdia noting size and reach considerations (Omdia). Include it where SSE/CASB integration is strategic.
- **Tenable** is credible where exposure management is the operating system. *CSO* describes Tenable One Cloud Exposure as joining DSPM with threat and identity/workload context across cloud, SaaS, and on-premises environments (CSO). Compare it with Wiz and one data-centric specialist.
- **OneTrust** is relevant where privacy operations and governance own the project, and Microsoft lists it as a Purview DSPM partner (Microsoft Learn). The security team should verify effective-access, activity, attack-path, and remediation depth rather than assuming privacy discovery equals DSPM.
- **Native cloud-provider services** can be good baselines and enforcement targets, but a multicloud DSPM purchase should demonstrate what it adds beyond them. Native services also provide a useful control for coverage and scan-cost comparisons.

Products absent here may still be suitable for a narrow source or regional requirement. The market changes quickly; acquisitions since 2023 include Polar/IBM, Imperva/Thales, Laminar/Rubrik, Normalyze/Proofpoint, Dig/Palo Alto Networks, Eureka/Tenable, Securiti/Veeam, and Wiz/Google (CSO; Omdia). Vendor identity and roadmap stability are therefore evaluation criteria in their own right.

How to build the right shortlist

Scenario A: multicloud, cloud-native engineering estate

Start with **Cyera, Sentra, and Wiz**. Add **BigID** if source heterogeneity, privacy, or governance is substantial; add **Palo Alto Networks** if Cortex is strategic.

The comparison reveals whether the buyer values deeper independent data context or native cloud attack-path context. Require support for the exact managed databases, object stores, lakehouses, Kubernetes volumes, serverless services, code systems, vector databases, AI platforms, and cross-account identity constructs in use.

Scenario B: Microsoft 365, Copilot, file shares, and SaaS collaboration

Start with **Varonis, Microsoft Purview, and Cyera or BigID**. Add **Rubrik** when resilience and Copilot readiness are linked.

This is principally an effective-access and usage problem. Test nested groups, external sharing, link types, guests, stale identities, sensitivity labels, endpoint/browser movement, and Copilot reachability. Microsoft's new DSPM can integrate partner insights, so "Purview plus specialist" may be more realistic than "Purview versus specialist" (Microsoft Learn).

Scenario C: privacy, records, governance, and security convergence

Start with **BigID, Securiti/Veeam, and Cyera**. Add **OneTrust** if already strategic.

Test one shared inventory across privacy impact, retention, residency, owner, security exposure, and AI use. The winning product should reduce duplicate scans and reconciliation—not merely place several modules behind one login.

Scenario D: regulated database and cryptographic control

Start with **Thales, IBM Guardium, and BigID or Cyera**. Add the incumbent database-activity monitoring vendor.

Make the vendors show discovery through protection: locate regulated fields, resolve access, identify a control violation, tokenize/mask/encrypt or change access, preserve evidence, and verify that the application still works.

Scenario E: data exfiltration, insider risk, and DLP modernization

Start with **Varonis, Proofpoint, and Cyera**; consider **Concentric AI** for unstructured semantic DLP and **Microsoft Purview** for a Microsoft-centered endpoint/productivity estate.

DSPM alone is primarily posture. Require the combined solution to observe movement, join it to user/entity behavior and content context, enforce policy, support investigation, and record the outcome.

Scenario F: recovery and cyber resilience first

Start with **Rubrik and Securiti/Veeam**, then add **Cyera, BigID, or Varonis** as an independent discovery benchmark.

The decisive question is whether the platform covers the live production estate and effective access, not merely the copies it protects. Conversely, make the specialist demonstrate how critical discovered data is protected and recoverable.

A rigorous proof-of-value method

1. Define the data estate before inviting vendors

Create a source ledger with:

- platform and service;
- account, tenant, subscription, project, region, and network boundary;
- structured, semi-structured, unstructured, image, audio/video, code, vector, and model artifacts;
- production, development, analytics, backup, archive, and collaboration copies;
- identity and permissions model;
- activity/audit source;
- business owner and regulatory classes;
- current catalog, labels, DLP, encryption, retention, and recovery controls.

Do not let a vendor's easiest connectors define the test. *CSO* warns that nominal support for a cloud provider does not mean support for all its data services and recommends documenting exact coverage and roadmap (CSO).

2. Build a lawful labeled test corpus

Use synthetic or properly authorized internal data; do not introduce real regulated data merely to test a tool. Include:

- exact identifiers and valid checksums;
- invalid and near-match identifiers to expose false positives;
- organization-specific project names, contract language, source code, secrets, and intellectual-property markers;
- multilingual content;
- data in tables, free text, nested JSON, archives, scanned images, and common office formats;
- duplicates and transformed copies;
- deliberately stale, overexposed, encrypted, mislabeled, and unowned assets.

Record ground truth before scanning. Calculate:

- **Precision** = true positives / all positive findings.
- **Recall** = true positives / all actual positives.
- **False-positive rate** and **false-negative count** by class and source.
- **Time to first useful finding**, **time to stable inventory**, and **freshness after a change**.

Do not accept a vendor's aggregate "accuracy." Require per-class results, sampling assumptions, thresholds, and error analysis. Confidence scores should be calibrated: among findings scored at 90%, approximately 90% should be correct on representative data.

3. Test the whole access graph

For each sensitive asset, compare:

- direct and inherited grants;
- nested groups;
- role and policy conditions;
- cross-account or cross-tenant trust;
- public and anonymous links;
- external guests;
- service principals, workload identities, API tokens, and AI agents;
- dormant identities and stale privileges;
- network reachability and infrastructure exposure;
- actual access and anomalous use.

The expected output is not a raw permission list. It is an explainable path: **identity** → **entitlement** → **reachable workload or data service** → **sensitive object**, with evidence and a safe corrective action.

4. Measure scan architecture and cost

"Agentless," "in place," and "data never leaves" are not interchangeable.

Require a data-flow diagram and packet-level or log-based verification showing:

- control plane and scanner locations;
- APIs, snapshots, clones, queries, crawlers, and agents used;
- raw content, samples, hashes, embeddings, labels, metadata, and findings transmitted;
- encryption, keys, tenancy, retention, subprocessors, support access, and deletion;
- required IAM privileges and separation of duties;
- impact on production databases and cloud bills;
- incremental-scan method and cadence;
- failure, throttling, retry, upgrade, and disaster-recovery behavior.

Rubrik's own buyer guidance, for example, argues that scanning should occur inside the customer environment and avoid redundant rescans (Rubrik). That is a sensible test criterion, not proof that every deployment achieves it.

5. Score closed-loop remediation

For ten representative findings, measure:

- whether the recommendation names the exact risky path and affected owner;
- whether it proposes the least disruptive fix;
- whether business and security owners can approve;
- whether exceptions expire;
- whether the action is native, API-driven, SOAR-driven, or ticket-only;
- whether rollback is possible;
- whether the product verifies closure at the source;
- whether the full evidence trail reaches SIEM, ITSM, GRC, and audit systems.

A dashboard finding is not a risk reduction. The outcome metric is verified reduction in reachable sensitive data, stale privilege, public exposure, unprotected copies, or policy violations.

6. Run blind, paired comparisons

Give each vendor the same sources, time window, privileges, corpus, and success criteria. Do not allow one to receive broader access or extra tuning without recording it. Have evaluators review findings without vendor labels where practical. Preserve raw exports so the team can compare duplicates, unique true positives, unique false positives, and missing assets.

7. Evaluate operations for 30–60 days

Initial scans flatter tools. The real test is change:

- discover a new store;
- copy sensitive content;
- alter a permission;
- create an external link;
- change a label;
- rotate an identity;
- add an AI data source;
- remediate and attempt recurrence.

Track alert latency, drift, reopened issues, analyst minutes per finding, owner response, connector failures, scan spend, and API throttling. A tool that finds 10% more theoretical issues but doubles triage work may be the worse control.

Suggested weighted scorecard

Weights should change by scenario, but this is a defensible default:

Criterion	Weight	Required evidence
Verified source and data-type coverage	15%	Successful scan of buyer-selected difficult sources
Classification precision, recall, and explainability	15%	Ground-truth corpus and error analysis
Effective access and identity context	12%	Verified complex permission paths
Risk prioritization and attack-path context	8%	Findings reproduced and ranked against known risk
Activity, lineage, and DDR	8%	Timed access/movement tests
Remediation and verification	12%	Closed-loop changes with approval and rollback
Architecture, privacy, and security	10%	Data-flow validation, least privilege, residency, retention
Performance and total cost	8%	Scan impact, cloud cost, licensing and labor model

Criterion	Weight	Required evidence
Integrations and operating workflow	5%	Working SIEM, ITSM, SOAR, catalog, IAM, DLP flows
Administration, reliability, and support	4%	30–60 day operational record and references
Vendor/roadmap and commercial resilience	3%	Contractual commitments, exportability, exit plan

Apply pass/fail gates before weighted scoring:

- no uncontrolled extraction of prohibited content;
- no unacceptable production impact;
- required residency and regulatory terms;
- complete support for the most critical sources;
- usable export/API and deletion on termination;
- no critical remediation without approval, audit, and rollback.

Commercial and contracting diligence

Public pricing is thin and variable. *CSO* reported in May 2026 that products are expensive, some charge by terabyte, marketplace subscriptions may start around $30,000 for small networks, and buyers should plan for at least $100,000 annually, rising with data volume (CSO). These are market-guide estimates, not quotes; obtain binding proposals.

Normalize bids against:

- discovered, scanned, sampled, and rescanned terabytes;
- files, objects, rows, records, repositories, accounts, users, or identities;
- production versus backup copies;
- connectors and premium sources;
- scanners, regions, business units, tenants, and cloud accounts;
- modules for DSPM, DDR, DLP, AI security, privacy, governance, access governance, and remediation;
- data egress, snapshots, warehouse compute, API charges, logging, and SIEM ingestion;
- deployment services, custom classifiers, support, and managed operations.

Contract for source coverage and outcomes, not just product access. Include:

- a named source/service/version matrix;
- service levels for connector failures and critical classification defects;
- notice and remedies for retired connectors;
- geographic processing and subprocessor commitments;
- vulnerability disclosure, penetration testing, incident notice, and secure-development evidence;
- customer ownership of policies, labels, findings, and exports;
- usable bulk export and deletion at termination;
- price protection for data growth;
- acquisition/change-of-control protections where material;
- acceptance criteria tied to the proof of value.

Practical red flags

- **A logo wall without a field-level support matrix.** A connector may inventory a tenant but not inspect content, effective permissions, or activity.
- **Precision without recall.** High precision can coexist with dangerous misses.
- **"No data leaves" without a data-flow diagram.** Metadata, samples, embeddings, screenshots, logs, or support bundles may still cross boundaries.

- **"Real time" based on daily scans.** Separate posture cadence from event-driven activity monitoring.
- **A risk score that cannot explain its path.** Analysts need evidence, not numerology.
- **Remediation that only opens tickets.** Ticketing is useful, but it is not enforcement or verified closure.
- **Automatic remediation without blast-radius analysis and rollback.** Removing access can break production and legal holds.
- **AI claims without model boundaries, evaluation data, and human review.** Microsoft itself warns that generative investigation can be inaccurate or incomplete (Microsoft Learn).
- **Roadmap counted as shipping capability.** This is especially important after acquisition and re-platforming.
- **A suite discount that avoids a specialist benchmark.** Consolidation is valuable only if the embedded DSPM meets the hard requirements.
- **Compliance reports presented as compliance.** A tool can produce evidence; accountability and control effectiveness remain organizational.
- **No exit path.** Findings, inventories, owners, labels, exceptions, and history must be exportable.

What the evidence does—and does not—support

The evidence supports a robust conclusion that the shortlisted products are credible enough to test and that their centers of gravity differ. Omdia's assessment supports BigID, IBM, OpenText, and Thales as 2025 leaders and Concentric AI, Rubrik, Securiti, Sentra, and Varonis as challengers; *CSO*'s 2026 review independently identifies Cyera, Microsoft, Palo Alto Networks, Proofpoint, Sentra, Tenable, Thales, Varonis, Veeam, and Wiz for investigation (Omdia; CSO). The overlap and differences show why no single analyst list should determine the RFP.

The evidence does **not** establish:

- which vendor has the highest classification recall on the buyer's data;
- which calculates effective access most accurately in the buyer's identity model;
- which is cheapest at the buyer's data volume and scan cadence;
- which creates the fewest operationally irrelevant alerts;
- whether advertised scale or time-to-value generalizes;
- whether acquisition-led integrations are complete in a specific tenant and region;
- whether AI-generated classifications or remediation suggestions are reliably correct;
- whether using one suite reduces total operating cost after all modules and labor.

Those are empirical questions for the proof of value. Public customer stories can demonstrate plausibility, not general efficacy. For example, Rubrik documents a fintech customer discovering and securing 7 PB and reporting 90% more visibility with its Laminar-derived DSPM (Rubrik customer story). Because this is vendor-published, unnamed-customer evidence without a controlled comparator or disclosed measurement method, it should inspire a scale test—not substitute for one.

Final recommendation

For most large enterprises, issue a focused proof-of-value invitation to **Cyera, BigID, and either Varonis or Wiz**, chosen according to whether collaboration/permissions or cloud attack paths dominate. Then:

- substitute **Sentra** for one candidate when the environment is cloud-data-platform heavy and the team wants a focused specialist;
- add **Microsoft Purview** for Microsoft- and Copilot-centric estates;
- add **Securiti/Veeam** for privacy/governance/AI/resilience convergence;
- add **Proofpoint** for DLP, human risk, lineage, and insider/exfiltration use cases;
- add **Thales** for cryptographic protection;
- add **Rubrik** for recovery-led data security;
- add **Palo Alto Networks** for Cortex consolidation;
- add **Concentric AI** for semantic unstructured-data control;
- include **IBM Guardium** for incumbent and regulated-database programs.

The winner should not be the product that discovers the most objects or produces the most severe dashboard. It should be the one that, on representative data and with acceptable privileges and cost, most accurately finds the sensitive assets that matter, explains who or what can reach them, detects meaningful change, and safely closes exposure in the organization's existing operating model.

Visual assets

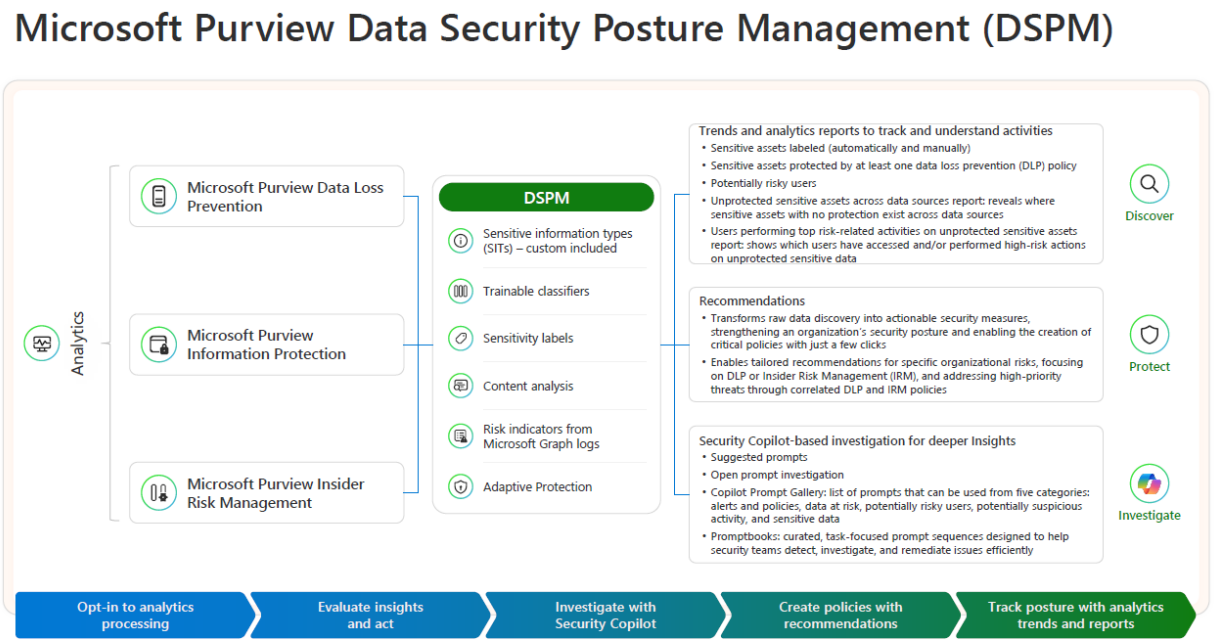

Figure 1. Microsoft Purview DSPM deployment model. Source credit: Microsoft Learn, "Deploy and use Data Security Posture Management". This is a vendor-authored workflow visual, not an independent comparison.

No additional source visuals were included. The independent Omdia report contains useful comparative figures, but its PDF states that unauthorized reproduction is prohibited; vendor marketing hero images and decorative stock imagery would not add evidentiary value.

Sources

Claim supported	Evidence type	Source	Direct URL
DSPM products span five use cases; coverage and integration vary; exact source support, architecture, cadence, and pricing require diligence; ten current vendors are worth investigation	Independent trade-publication buyer's guide drawing on analyst research and vendor inquiries	*CSO*, "DSPM buyer's guide: Top 10 data security posture management tools," 27 May 2026	https://www.csoonline.com/article/2075321/top-12-data-security-posture-management-tools.html
2025 comparative positioning; most products retain gaps; BigID/IBM/OpenText/Thales leaders; Concentric/Rubrik/Securiti/Sentra/Varonis challengers; Thales strengths and limitations	Independent analyst comparative assessment; hosted by a vendor, copyright Omdia	Omdia Universe: Data Security Posture Management, 2025	https://cpl.thalesgroup.com/sites/default/files/content/analyst_research/Omdia-Universe-Data-Security-Posture-Management-DSPM-2025-report.pdf
Cyera's stated source coverage, agentless architecture, classification, scale, precision, customer metrics, and remediation actions	Vendor product documentation and vendor claims	Cyera DSPM	https://www.cyera.com/platform/dspm
BigID's stated discovery/remediation model, classifiers, source breadth, privacy/governance/AI context, and performance claims	Vendor product documentation and vendor claims	BigID DSPM	https://bigid.com/data-security-posture-management/
BigID's stated 500+ native data sources, data-type coverage, deployment breadth, and integrations	Vendor product documentation and vendor claims	BigID DSPM + AI-SPM	https://home.bigid.com/dspm-demo
Varonis's stated SaaS/IaaS/file coverage, third-party app controls, activity trail, behavior analytics, and intervention	Vendor product documentation	Varonis SSPM/data security	https://www.varonis.com/solutions/sspm
Sentra's stated in-place scanning and cloud/SaaS/on-premises scope	Vendor product documentation	Sentra DSPM glossary	https://sentra.io/glossary/data-security-posture-management
Wiz's stated cloud/code/AI discovery, access governance, compliance, and cloud-graph risk correlation	Vendor product documentation	Wiz DSPM	https://www.wiz.io/solutions/dspm
Wiz's extension of DSPM classifiers to code, pull requests, IDEs, and CI/CD	Vendor technical announcement	Wiz, "Wiz extends DSPM capabilities to code bases," 12 March 2024	https://www.wiz.io/blog/wiz-extends-dspm-capabilities-to-code
New Microsoft Purview DSPM functions, Microsoft and third-party coverage, partner integrations, objectives, AI observability, remediation, and current limitations/previews	Primary technical documentation	Microsoft Learn, "Learn about Data Security Posture Management"	https://learn.microsoft.com/en-us/purview/data-security-posture-management-learn-about
Microsoft's warning that generative-AI results may be inaccurate or incomplete	Primary technical documentation	Microsoft Learn, Data Security Posture agent	https://learn.microsoft.com/en-us/purview/data-security-investigations-posture-agent
Microsoft DSPM deployment workflow and reproduced visual	Primary technical documentation	Microsoft Learn, "Deploy and use Data Security Posture Management"	https://learn.microsoft.com/en-us/purview/deploymentmodels/depmod-dspm-intro

Claim supported	Evidence type	Source	Direct URL
Veeam completed its $1.725bn Securiti acquisition on 11 December 2025 and states the combined scope	Primary corporate announcement	Veeam acquisition completion	https://www.veeam.com/pt/company/press-release/veeam-acquires-securiti-ai.html
Proofpoint's current Normalyze-derived DSPM positioning and claimed cloud/on-premises classification, access-risk prioritization, and remediation workflows	Vendor product documentation	Proofpoint, "Normalyze is now Proofpoint"	https://www.proofpoint.com/us/normalyze-is-now-proofpoint
Normalyze/Proofpoint's stated one-pass scanning, access graphs, Risk Navigator, data valuation, customization, and AI functions	Vendor technical/product guide	Proofpoint DSPM buyer guide	https://www.proofpoint.com/us/blog/dspm/an-essential-guide-to-data-security-posture-management-dspm-solutions
Normalyze's stated data-in-motion discovery and lineage	Vendor technical announcement	Proofpoint/Normalyze, data in motion and lineage	https://www.proofpoint.com/us/blog/dspm/normalyze-data-in-motion-and-data-lineage
Palo Alto Networks' current agentless, Dig-derived DSPM module and tenant deployment note	Vendor field guide	Palo Alto Networks Prisma Cloud Field Guide	https://start.paloaltonetworks.com/rs/531-OCS-018/images/Prisma%20Cloud%20Field%20Guide_v2.pdf?version=2
Rubrik's stated mature-DSPM criteria, in-environment scanning principle, risk/access/governance model, and scan-efficiency guidance	Vendor product guide	Rubrik, "What Is DSPM?"	https://www.rubrik.com/insights/what-is-dspm
Rubrik's stated Microsoft 365 Copilot discovery, permission remediation, and Purview label functions	Vendor data sheet	Rubrik DSPM for Microsoft 365 Copilot	https://www.rubrik.com/content/dam/rubrik/en/resources/data-sheet/ds-dspm-for-m365-copilot.pdf
Documented 7 PB/90%-visibility fintech example, subject to vendor-case-study limitations	Vendor customer story	Rubrik fintech customer story	https://www.rubrik.com/customers/fintech
Concentric AI's stated semantic, context-aware coverage of data at rest and in motion	Vendor product documentation	Concentric AI	https://concentric.ai/