



The New Standard for Data Access Governance: Reimagined for the Age of AI

EXECUTIVE SUMMARY

Most organizations have governance tools that can map permissions and generate entitlement reports. What they can't do is tell you whether the data behind those permissions actually matters, whether it's being used, or whether an AI agent accessed it at 3am. True DAG requires knowing what the data contains, who and what can reach it, and what's actually happening — without hindering the business.

Why organizations are rethinking DAG — and what the next generation looks like

The Limits of Permissions-Centric Governance

Traditional data access governance tools were built around a single question: who has access to what? That was a reasonable starting point. But it turns out, that question alone doesn't tell you much.

Knowing that a service account has read access to 200 tables is not the same as knowing that those 200 tables contain Social Security numbers, financial records, and health data. Knowing that a user has broad access to a file share doesn't tell you whether they've accessed anything, or whether the files they can reach are sensitive at all. And knowing which permissions exist today tells you nothing about whether they've ever been used — or whether a new AI agent was just granted the same access last Tuesday.

Governance tools that treat access as a permissions problem — without a data intelligence layer underneath — produce entitlement reviews that are technically complete and operationally meaningless. Teams spend cycles certifying access to things that don't matter while genuinely high-risk access sits unaddressed. Oversharing accumulates quietly. Permissions granted for a project years ago remain in place long after the project ends.

If your governance layer and your data classification layer are different systems, your governance decisions are always one step removed from the data risk they're supposed to be managing.

The Expanding Access Risk Surface

The governance surface has expanded dramatically in the past two years, and it's not slowing down.

The traditional challenge — over-permissioned employees, dormant service accounts, nested group memberships that no one fully understands, external collaborators with access that was never formally reviewed — hasn't gone away. But layered on top is an entirely new category of access risk: AI agents, MCP-connected tools, LLMs with data access, and automated pipelines that operate continuously, at scale, and often with permissions scoped far too broadly for what they actually need to do.

These actors don't behave like humans. They don't browse; they query. They don't make mistakes the way humans do; they execute against whatever permissions they've been granted. Behavioral baselines built on human access patterns won't catch them. Governance workflows built around quarterly access reviews won't govern them. And access reports that don't distinguish between human users, external collaborators, and autonomous agents will miss the risks that matter most.

What Integrated DAG Actually Requires

Data access governance that can answer the questions that matter — not just "who has access" but "who has access to what sensitive data, has that access been used, and is it appropriate given what the data contains" — requires three capabilities working as one:

Data Security Posture Management

Provides the intelligence layer. DSPM continuously discovers, classifies, and catalogs sensitive data across structured databases, unstructured file stores, cloud environments, and data pipelines. Without this, access governance is operating blind — making decisions about permissions without understanding what those permissions expose.

Data Access Governance

Provides the control layer. BigID delivers comprehensive visibility into all users, groups, service accounts, external collaborators, and AI agents with access to data — and automatically identifies overexposed sensitive information so teams can remove unnecessary permissions before they result in incidents.

Entitlement reviews are delegated and auditable: data owners can review direct permissions for specific files and choose to keep or revoke them, with decisions submitted directly or routed through ticketing systems like Jira and ServiceNow. Enforcement is proportional to actual data sensitivity, not just permission breadth.

Data Activity Monitoring

Provides the behavioral layer. BigID monitors user and system behavior across platforms — including Google Drive, AWS S3, Microsoft SharePoint, OneDrive, and NetApp — tying activity directly to sensitive data access. This context lets analysts identify actual risks based on real behavior rather than assumptions.

Out-of-the-box threat detection policies correlate activity with classification metadata in near real time, surfacing detailed context — file paths, sensitivity labels, access patterns — for security alerts like external users accessing sensitive data. The result is fewer false positives and faster, more confident response.

When these three capabilities share a data model, governance decisions become grounded in reality. Entitlement reviews surface access that actually matters. Anomaly detection is accurate. Remediation is proportional.

Governance That Extends to AI

For organizations deploying AI agents and agentic workflows, integrated DAG isn't optional — it's the prerequisite for deploying AI responsibly.

Every AI agent, every MCP-connected tool, every automated pipeline that touches data needs to be governed the same way human access is governed: with discovery-backed sensitivity context, access entitlement visibility, and behavioral monitoring that can detect when something unexpected is happening.

This means labeling and classifying what data AI systems can access. It means scoping AI access rights based on what the data actually contains. It means monitoring AI data access continuously, with the same rigor applied to high-risk human access. And it means being able to act — revoke access, trigger workflows, enforce policies — when something looks wrong.

Organizations that treat AI governance as a separate problem from data access governance will find themselves running parallel systems that can't talk to each other. The smarter approach is to extend the same governance model that applies to humans and service accounts to cover AI agents natively — from the same platform, against the same data intelligence foundation.

Operating at Scale, Without Slowing Down the Business

One reason governance programs stall is that they create friction. Blanket access revocations break workflows. Overly aggressive entitlement reviews generate noise and resentment. Security teams end up manually triaging backlogs of access violations without enough context to know which ones actually matter.

BigID addresses this directly. A unified security overview dashboard gives teams a consolidated view of their data security posture — highlighting the risk drivers and open access violations that warrant attention, rather than surfacing everything at once. Delegated review workflows distribute the work to data owners who have the context to make good decisions. And MCP server and API integration means internal tools and third-party systems can query catalog findings and orchestrate governance workflows programmatically — reducing manual effort and enabling automated, interconnected governance at scale.

The goal isn't to lock everything down. It's to ensure that access to sensitive data — by humans, service accounts, and AI agents alike — is intentional, appropriate, and continuously monitored.

Conclusion

The era of permissions-only governance is over. Access risk today spans employees, external collaborators, service accounts, and AI agents operating at machine speed. The data those identities can reach is spread across cloud environments, databases, file stores, and AI pipelines. And the standard has shifted — from "who has access" to "who has access to what sensitive data, has that access been used, and is it appropriate."

BigID delivers data access governance built on data intelligence: integrated DSPM, DAG, and DAM in a single platform, purpose-built for the access risk surface organizations face today — and the agentic one they're building for tomorrow.

CUSTOMER VALIDATION

Proof in Action

“BigID gives me better visibility into sensitive data, helps prioritize security protections, reduces attack surfaces, strengthens compliance, and increases operational efficiency overall — a strategic pillar of our AI-First Cybersecurity Transformation”

- CISO, Global Healthcare Company

“We needed to automate processes and data management across systems for the data of a few million customers across a lot of systems and data. BigID was the one solution that did this in the most efficient and sophisticated way - and had more use cases we could add on moving forward.”

- Chief Privacy Officer, Global Telecoms Company

ANALYST RECOGNITION

Validated by Industry Leaders

- 2025 Company of the Year for AI Governance - Frost & Sullivan
- Named a Leader across all DSPM research including Omdia DSPM Universe, CB Insights for DSPM, TAG, GigaOm, Frost & Sullivan
- Named a Leader in all Privacy Management evals including Forrester Privacy Management Wave & IDC Privacy Compliance MarketScape
- BigID named a Leader across 4 GigaOm evals: DSPM, Data Security Platforms, Unstructured Data Management, Data Access Governance
- Named in Gartner's 2025 DLP Market Guide
- Named in Gartner's 2025 AI TRiSM Market Guide
- Represented 30+ times in Gartner's 2025 Hype Cycles for Security, Privacy, AISPM

FREE RISK ASSESSMENT

Discover Your Data & AI Risk

An agentless, cloud-native assessment to enable security, privacy, and compliance: pinpoint where sensitive data and AI risks live - and how to fix them.

What You Get:

- Complementary 2 week engagement on real data to gain visibility on risk across data and AI
- Executive-ready report identifying high-risk data and AI across cloud environments with next-step recommendations

NEXT STEPS

Ready to Connect the Dots in Data & AI?

Schedule your free risk assessment or ask for a deeper demo across DSPM, DLP, AI SPM, Privacy Automation, AI Risk, and Data Lifecycle Management

▶ www.bigid.com/demo

▶ www.bigid.com/risk-assessment

▶ info@bigid.com



BigID Next

"The features they offer are the moat that they have built around the product. This is what makes them the market leader in Data Privacy and Protection."



IT Security & Risk Management Associate - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.



BigID Next

"Exceptionally powerful and easy to use tool that integrated smoothly into our environment and delivered value from day one. The BigID team is fabulous to work with."



Senior Vice President & Chief Data Officer - Banking

Gartner
Peer Insights.

Gartner® and Peer Insights™ are trademarks of Gartner, Inc. and/or its affiliates. All rights reserved. Gartner Peer Insights content consists of the opinions of individual end users based on their own experiences, and should not be construed as statements of fact, nor do they represent the views of Gartner or its affiliates. Gartner does not endorse any vendor, product or service depicted in this content nor makes any warranties, expressed or implied, with respect to this content, about its accuracy or completeness, including any warranties of merchantability or fitness for a particular purpose.