

The Data Control Plane for Responsible AI



EXECUTIVE BRIEF

AI agents operate on your data. What they can access, what they can modify, and whether any of it is governed determines whether your AI is responsible in practice — not just in principle. BigID is the data control plane that makes responsible AI operational and trust in AI systems earned, not assumed.

Responsible AI Has an Infrastructure Problem

Responsible AI is one of the most cited priorities in enterprise technology. It is also one of the least operationalized. Most organizations have policies, principles, and governance frameworks. What they are missing is the infrastructure to enforce any of it — and the visibility to know when something goes wrong.

AI agents do not make decisions in isolation. They make decisions based on data — data they retrieve, data they modify, data they carry from one action to the next. The quality, classification, and governance of that data determines whether the agent behaves responsibly. No alignment framework, bias audit, or governance policy changes that equation. Only controlling the data does.

Forrester's Responsible AI Solutions Landscape, Q2 2026 identifies the critical gap the market has not closed: most organizations still rely on point-in-time, reactive RAI solutions built for static models, not autonomous agents operating across interconnected systems. The top disruptor Forrester calls out is the ability to observe and remedy agent behavior across multisystem decision chains — continuously, in real time. That is a data infrastructure problem. BigID is the answer.

The Data Control Plane for Agentic AI

Every production system has a control plane — the layer that governs what is allowed, enforces policy, and provides visibility into what is actually happening. For networks, it routes traffic and enforces rules. For cloud infrastructure, it manages resource allocation and access. For AI agents operating across enterprise environments, most organizations have built nothing equivalent.

BigID is the data control plane for agentic AI: the system that defines what data agents can access, enforces the policies that govern their behavior, and delivers continuous observability across every action they take. It is not a governance dashboard layered on top of your AI stack. It is the infrastructure running underneath it — the foundation that makes everything above it auditable, defensible, and worthy of organizational trust.

Three capabilities define what this means in practice:

1. Know what data your agents are operating on

Agents are only as reliable as the data context they operate within. If underlying data is misclassified, overexposed, stale, or ungoverned, agents will act on it regardless — producing outputs that are wrong, risky, or non-compliant. BigID continuously discovers and classifies data across cloud, on-prem, and hybrid environments. It knows what is sensitive, where it lives, who owns it, and what systems and identities can reach it. That is the baseline. Without it, responsible AI is a declaration, not a capability, and trust in AI outputs cannot be verified.

2. Control what agents are allowed to do

Access governance for AI agents is not optional — it is the prerequisite for responsible deployment at scale. When an agent can query customer databases, retrieve regulated records, or modify data across systems, the policies governing that behavior need to be enforced at the data level, not just documented somewhere. BigID applies the same identity and access governance it uses for human users to AI agents and agentic systems: defining least-privilege access, enforcing controls at the data layer, and detecting violations in real time. This is what separates responsible AI in practice from responsible AI on paper.

3. Maintain full observability across the AI lifecycle

You cannot have accountability without a record of what happened. BigID traces data from source through AI consumption, modification, and output — providing the lineage required for explainability, the audit trail required for regulatory accountability, and the continuous monitoring required to catch drift, anomalies, and policy violations before they become incidents. When an agent's behavior changes, BigID sees it. When a boundary is crossed, BigID acts. This is how organizations move from periodic audits to genuine, ongoing AI oversight.

The Missing Layer in Agentic Infrastructure

When organizations build agentic AI platforms, the conversation centers on models, orchestration frameworks, tool integrations, and interfaces. What consistently gets underbuilt is the data infrastructure those agents depend on. The assumption is that the data environment is already handled — that existing governance, access controls, and

classification systems are sufficient. For most enterprises, they are not.

Agentic AI introduces a fundamentally different risk profile from previous AI deployments. Traditional AI models produce outputs that humans review. Agents take actions — modifying data, triggering workflows, making decisions that propagate across systems — often with minimal human review in the loop. The data they access, the permissions they carry, and the trail they leave behind become the primary record of what happened and whether it was appropriate.

BigID addresses this gap as a core component of agentic infrastructure. It answers the questions no model can answer about itself:

- What data am I working with, and is it accurately classified and properly governed?
- Am I authorized to access this data under current policy and regulatory requirements?
- Am I operating within the boundaries this organization has defined?
- If something goes wrong, what did I touch, what changed, and when?

Organizations that deploy agentic platforms on top of BigID's data intelligence foundation do not just have AI that works. They have AI they can verify, audit, and stand behind — with regulators, with boards, and with customers. That is what earns durable trust in AI systems. Not a declaration, but a demonstrable foundation.

Why BigID

Most responsible AI solutions fall into two categories: governance and risk management tools that operate above the data layer, and model performance tools that operate inside the AI stack. Both categories assume that the data environment is already known, classified, and controlled. BigID is the only solution in this market built from the ground up on data intelligence — the ability to find, understand, and govern data at enterprise scale across every environment where it lives.

Every responsible AI requirement depends on that foundation:

- Explainability requires data lineage. BigID provides it.
- Alignment requires access control. BigID enforces it.
- Accountability requires continuous observability. BigID delivers it.
- Human oversight requires knowing what agents actually touched. BigID tracks it.

BigID does not replace governance and compliance layers. It makes them functional. Where most responsible AI solutions deliver periodic snapshots, BigID operates continuously — governing the data environment, enforcing policy, and surfacing changes as the AI stack runs. Forrester recognized BigID in its Responsible AI Solutions Landscape, Q2 2026, citing its focus on AI observability and lineage, policy management and

enforcement, and human oversight: the three capabilities that define a data control plane.

What This Means for Your Organization

For the CISO: Data Control Is Agent Control

AI agents are the new attack surface — and they arrive pre-authorized to access your most sensitive data. Overprivileged agents operating on unclassified or overexposed data create risk that no SIEM, DLP, or endpoint tool is designed to catch. The access governance problem for agents is structurally identical to the access governance problem for human identities: you need to know what they can reach, enforce what they should reach, and detect when they deviate.

BigID governs agent access to data with the same rigor it applies to human identity — least privilege by default, continuous monitoring, and real-time policy enforcement. For security leaders, responsible AI is not a model alignment question. It is a data control question. BigID is built to answer it.

For the CIO: Trustworthy AI Requires Infrastructure, Not Just Intent

AI investments fail when the data beneath them is wrong. Agents built on ungoverned, misclassified, or overexposed data produce unreliable outputs, create regulatory exposure, and erode the organizational confidence needed for AI adoption to scale. The harder problem is not building agents that work in a demo — it is building agents that work reliably, comply with policy, and can be defended when something goes wrong.

BigID is the infrastructure investment that makes AI deployments defensible and scalable. Not because it constrains what agents can do, but because it ensures they operate on the right data, under the right policies, with full visibility into every action. That is what makes AI trustworthy in the eyes of boards, auditors, regulators, and customers — and it starts with the data layer.

Responsible AI is not a model capability. It is a data infrastructure requirement. BigID is the data control plane that makes it real — governing what agents can access, enforcing the policies that define acceptable behavior, and delivering the observability that makes AI systems auditable and trustworthy by design.

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.