

Data Filtering for Confidential AI

Secuvy ensures only appropriate data ever enters your AI pipeline – filtered, governed, optimized.

Every Fortune 500 leadership team surfaces three key AI commitments: accelerate decisions, drive efficiency, create revenue. The budgets are real. The ambition is clear. What rarely gets asked is the harder question – what data are we actually going to feed it? Most enterprises are running on decades of accumulated, ungoverned data never designed for machine consumption. Sensitive material, duplicates, and unclassified records sit invisibly in the pipeline. The model can't distinguish appropriate from inappropriate – because no one has told it what's there.


The consequences land in three places: **Risk** – PII, PHI, IP, and classified data surfaced at inference, with audit obligations that don't care how it got there; **Cost** – GPU and Tier-0 storage burned on low-quality data, producing underwhelming results the CFO is already questioning; and **Time** – up to 80% of AI effort consumed by data preparation before a single workload runs. From months of engineering before a model goes live, to GPU cycles burned on the wrong data – the investment in enterprise AI is visible on every level; the return isn't.

"58% of IT leaders cite classifying data for AI as their top technical challenge."


– Komprise, 2026 State of Unstructured Data Management (5th annual survey)

The Solution: Secuvy – The fastest, most accurate data filtering solution on the market

Secuvy is the AI-native data filtering platform enterprises trust to govern what their models consume. Built for real-world complexity – not static rules – it continuously discovers, classifies, and governs data so only appropriate data reaches your AI pipeline. The result is complete visibility into what's in your data – and confidence that every model is fueled by sources that are filtered, appropriate, and fully accounted for. Every model is fueled by sources that are governed, fit for purpose, and fully accounted for – moving enterprises from reactive data management to continuous governance, ready for the next audit, the next earnings call, and the next AI initiative.



How Secuvy Works



Secuvy dashboard showing classification and business tagging mapped to data sources

Secuvy built the most accurate, self-learning AI model for continuous data classification and data recommendation — the foundation for AI data management, compliance and governance.

- **Self-Learning Filtering.** Secuvy continuously discovers and auto-classifies data using an unsupervised AI engine that adapts to real content and context. Accuracy improves as more data is processed — cutting false positives and eliminating constant rule management.
- **Non-Pattern Data Classification.** Secuvy does not rely on pattern matching. Anything can be classified: define the need, and Secuvy finds it in the data.
- **Contextual Linkage.** When related data shares sensitive signatures across systems, Secuvy connects those records and visualizes the linkage automatically — providing context that traditional classification tools simply can't see. Users, systems, access paths, and governing policies are all correlated into a clear, risk-weighted view of exposure.
- **Sensitivity-Driven Risk Visibility.** Rather than inferring risk from access paths alone, Secuvy anchors risk analysis in data sensitivity itself. Exposure stays visible even as data moves across platforms, users, and AI systems.
- **AI-Aware Governance.** Secuvy treats AI as a data access layer that must be governed using real data context. Guardrails are enforced based on what's actually sensitive — not generic assumptions or static policies.
- **Fast Filtering.** The fastest data filtering and classification engine on the market. Operating at the file-system level, it delivers the fastest path to classification — avoiding the pitfalls of protocol-level approaches.

Traditional solutions were built for a different era. Static classification requires manual tuning, and pattern-matching approaches with partial visibility routinely miss sensitive data. The result is excessive noise — and a loss of trust that anything is being classified correctly.

Solution Benefits

Lower Risk

Secuvy ensures you know exactly what's going into your AI engine. Your CISO and legal team can prove that PII, PHI, IP, and classified data is blocked from training and inference – with the lineage to defend in front of a regulator, an auditor, or a board.

Confidence to Decide

A verifiable data bill of materials becomes the system of record for every AI workflow. AI stops being an expensive experiment and starts being infrastructure – with the proof to put on the next earnings call.

Faster Time to Value


Replace months of manual data scrubbing with continuous, automated filtering. Move from idea to first results in days, not quarters – with nothing surfacing in legal review weeks or months down the road.

Lower Cost with Better Data

Cleaner, smaller datasets shrink GPU and Tier-0 storage footprints, eliminate duplicates, and remove erroneous data. Your CFO sees AI investment driving business value – not burning compute on a dataset nobody can vouch for.

Built for the NVIDIA AI Data Platform

Secuvy works with the NVIDIA AI Data Platform (AIDP) architecture to produce a highly accurate data bill of materials (dBOM) – the AI equivalent of the software bill of materials that became table stakes for secure software development. With Secuvy in front of AIDP, every GPU cycle is spent on the right data: governed, classified, and fit for purpose. Leaders who reach that bar stop having defensive conversations about AI in the boardroom and start having offensive ones.



Secuvy dashboard showing the complete data bill of materials (dBOM) for an organization

“I can see everything with Secuvy, so I'm confident my data is ready.”

— Enterprise Data Architect

Secuvy: Filtering Data for Confident AI

Secuvy ensures every AI model is fueled by appropriate sources that are governed, fit for purpose, and fully accounted for — with the lineage to defend in front of a regulator, an auditor, or a board.

At it's core, Secuvy provides a verifiable data Bill of Materials (dBOM) as the system of record for every AI pipeline — proving that the right data, and only the right data, fueled every model. This gives enterprises what most AI programs are missing: the ability to scale without scaling the risk surface, convert AI investment into outcomes leadership can stand behind, and move from pilot to production without hitting the data preparation and compliance bottlenecks that stall most initiatives.

The result is an AI strategy that compounds returns — not one that stalls on risk, cost, and time.

Go to **secuvy.ai/demo** and schedule 30 minutes with our team to see your path to success.



**Know what's in your data
estate before your AI does.**

Schedule 30 minutes with our team to
see your path to success.



**[secuvy.ai/
demo](https://secuvy.ai/demo)**