


AI projects hold exciting promise for driving business value through insights and automation. But there's a snag that often goes unnoticed until it's far too late: AI data ingestion pipelines routinely pull in vast amounts of "junk" or irrelevant files, muddying the datasets and driving up costs. This "dark data" not [unstructured data management](#) only bloats storage and backup budgets but also creates serious security and compliance risks. In this post, we'll explore what dark data really is, why it accumulates, and strategies — including modern tools like Komprise Intelligent Data Management and DataStackHub solutions — to gain visibility, curate clean datasets, and exclude junk from AI pipelines.

What is Dark Data and Why Does It Accumulate?

The term dark data refers to information assets organizations collect, process, and store during regular business activities but fail to use for analytics, AI, or other purposes. One client recently told me they thought they could save money but ended up paying more. Gartner has identified dark data as a massive untapped reservoir—often comprising aging files, test datasets, multiple redundant copies, and more.

Why does dark data pile up? Think of your file storage like a cluttered attic stuffed with old boxes no one wants to sort. Over time:



- Files accumulate from legacy projects, end-of-life applications, and employees who leave.
- Repeated backups, copies, and replications create duplicates.
- Lack of tagging or metadata means nobody knows who owns or needs many files.
- Automation and AI projects suck in all available data by default — including irrelevant files — which just adds to the problem.

Dark data is often unstructured—spreadsheets, PDFs, images, videos, email archives—making it especially challenging to detect and manage without specialized tools that provide *unstructured data visibility and discovery*.

Why Dark Data is a Cost Driver for AI Data Ingestion

Dark data's impact hits your bottom line in multiple expensive ways:

Cost Category Description AI Impact Storage Paying for capacity to hold excessive, irrelevant files that aren't contributing to business outcomes. AI pipelines ingest this unnecessary bulk, inflating cloud storage or on-prem NAS bills. Backup and Replication Dark data multiplies backup windows and replication footprints, increasing operational overhead. Longer backups and replication increase latency and reduce agility for AI model refreshes. Disaster Recovery (DR) DR plans must account for all data volumes, even irrelevant files, impacting recovery times and costs. Ensuring AI repositories recover quickly becomes harder when bulked up with junk data. Security and Compliance Hidden PII, PHI, or other regulated info buried in dark data risk breaches, fines, and reputational damage. AI ingestion widens exposure to sensitive info that may lack proper controls or auditing.

According to Gartner, organizations can be paying up to 30% in extra costs for managing dark data alone — and AI projects exacerbate these expenses if data isn't properly curated.

Unstructured Data Visibility and Discovery: The First Step to Exclusion

Before you can exclude junk files from AI ingestion, you need to **know what your storage contains**. This is where unstructured data visibility and discovery come in. Tools like Komprise Analysis provide deep insights across file shares, NAS devices, cloud buckets, and object stores.

What Visibility Tools Offer:

- **Ownership Identification:** Find out who last modified or owns files — essential to ask “who owns this data?” before deciding its AI value.
- **Age and Access Patterns:** Detect stale data that nobody touches, often prime candidates for exclusion.
- **Data Type and Sensitivity:** Discover hidden sensitive info like PII or PHI embedded in unstructured files.
- **Duplicates and Redundancy:** Identify repeated copies inflating your footprint.

This knowledge vault enables data engineering and AI teams to build curated datasets that are relevant, sensitive-aware, and cost-efficient.

How to Build Curated Datasets and Exclude Dark Data from AI Pipelines

Once you achieve visibility, the real work begins: actively managing data to exclude junk from ingestion.

Strategies to Stop Junk Files from Creeping In:

1. **Implement Intelligent Data Tiering and Archiving:** Tools like Komprise Intelligent Data Management automate movement of cold, older files into cost-effective archive tiers or cloud storage designed for infrequent access.
2. **Apply Data Classification and Tagging:** Use content scanning and metadata enrichment to tag datasets with ownership, sensitivity, business value, and retention policies.
3. **Set Exclusion Filters in AI Ingestion Workflows:** Integrate with visibility tools' APIs to exclude files based on age, owner, size, or sensitivity before pulling into models.
4. **Maintain a “Mystery Shares” List:** Regularly track unknown or orphaned shares that nobody claims ownership of — commonly hotbeds of junk data — and remediate or quarantine as per legal guidance.

5. **Collaborate Across Teams:** Engage security, legal, and business owners early to ensure compliance policies are embedded in data prep.


Benefits of This Approach

- **Reduce Storage Costs:** Eliminate paying for space to store files with no AI value.
- **Speed Up Model Training:** Smaller, cleaner datasets cut down data loading and processing time.
- **Reduce Risk:** Avoid ingesting ungoverned PII or PHI that could cause data breach fines or compliance headaches.
- **Future-Proof Data Use:** Create a governance framework so AI initiatives don't compound dark data challenges for the next project.

Case Study: A Fortune 1000's Journey with Komprise

One Fortune 1000 company faced ballooning cloud storage bills as their AI team's ingestion workflow indiscriminately pulled every accessible file share. Using Komprise Analysis, they discovered that 40% of the files were old project archives or duplicates. Implementing Komprise Intelligent Data Management, they tiered and archived cold data and set ingestion filters excluding files untouched for over 2 years.

The results:



- 35% reduction in storage costs within 6 months
- Faster turnaround times on AI model training by focusing on active, business-critical data
- Improved compliance posture by proactively identifying sensitive files hidden in dark data

Key Takeaways

- **Dark data is the “junk in the attic” for your AI pipelines; it accumulates quietly but carries big costs and risks.**
- **Visibility into unstructured data is essential — tools like Komprise Analysis help identify the bad apples before ingestion.**

- **Curating datasets through tiering, classification, and exclusion filters can stop junk files from bloating costs or exposing sensitive info.**
- **Cross-team collaboration, especially involving security and legal, ensures governance is built into AI data prep.**
- **Avoid the common mistake of “just pull everything into the AI data lake” — smarter curation pays off in agility, cost savings, and risk avoidance.**

Final Thoughts

AI offers tremendous opportunities, but success depends on feeding it the right data, not garbage. If your ingestion workflows are pulling in junk files, take a step back and ask “who owns this data?” and “do we really need to include it?” Armed with unstructured data visibility and intelligent management platforms like Komprise, you can stop AI ingestion pipelines from hauling in wasted storage, security headaches, and irrelevant noise — and instead harness clean, curated datasets that power real business value.

For more insights on managing dark data and enabling AI-ready datasets, explore resources from Komprise, Gartner, and DataStackHub.