

Everpure Data Intelligence File Edition

Know what's in your storage before deciding what to do with it.

You know exactly how much data you're storing. But what that data is, how old it is, how much risk it poses, and whether it's stale are questions that capacity reporting can't answer.

Everpure Data Intelligence™ File Edition extends existing functionality to metadata across every file you already have—names, paths, sizes, owners, and dates—without accessing the file content itself. One scan turns that metadata into a score that tells you which repositories or folders to review first.

Everpure Data Intelligence File Edition answers those questions without opening or reading the file.

Why intelligence matters for files

A share showing terabytes of data and growing tells you nothing about what's inside. Most of it hasn't been opened in years, and a folder buried in it might be accessible by anyone holding the link. So storage optimization decisions get made on age alone, cleanup stalls while teams determine what's safe to remove, and when the AI team asks which data it can use, no one has a defensible answer.

The questions are shifting from how much to what and whose, and they land differently depending on your seat. If you run storage and infrastructure, you need reclaim targets, not another capacity number that keeps climbing. If you're the CISO, you need to know where exposure concentrates before an audit forces the question. If you lead data or AI initiatives, you need to know which datasets are current enough to trust with a model. Whatever your role, you need to discover, understand, and manage that data before you can act on it.

How File Edition works

Everpure Data Intelligence File Edition connects to the sources you already run (Amazon S3, OneDrive, SharePoint, Google Workspace, CIFS/SMB, and NFS) and collects metadata only (names, paths, sizes, owners, dates, and location). No file is opened at any point.



Hybrid and multicloud visibility

Connect to S3, OneDrive, SharePoint, Google Workspace, CIFS/SMB, and NFS without ever opening a file.



Scan metadata only, every time

Only look at metadata insights to inform infrastructure efficiency, security, and data decisions.



Surface hotspots that require deeper analysis

Pinpoint high-signal datasets for contextual classification scans.




FIGURE 1 File intelligence workflow

Five metrics run across the top of the file intelligence dashboard: total files scanned and the capacity they consume, the percentage concentrated in file types you care about, the percentage externally accessible, the percentage that falls outside your freshness window, and the number of repositories and folders flagged for review.

Eight metrics run across the dashboard in total. Four of them—file type density, staleness, external access, and entitlement concentration—are used to calculate the signal score.

The signal score delivers prioritized, actionable insights

Raw storage metrics are not comparable. A share at 40% stale and a bucket at 40% stale are not equally concerning if one is an outlier and the other is normal for your estate.

The signal score is a “look-here-first” metric. The higher it runs, the further that share, bucket, or folder sits from its peers, and the more it’s worth your attention.

File intelligence measures four signals (file type density, staleness, external access, and entitlement concentration) and converts each into a percentile rank against comparable peers. Each repository is ranked against every repository in your environment, and each folder is ranked against all folders in the same location. A ranking of 95 means the repository or folder requires more attention than 95% of its peers.

Those rankings blend into one score between 0 and 100. Half of the score comes from the worst individual signal, ensuring one severe problem is never averaged away. The other half comes from the weighted mix of all signals, so a source with multiple risk factors ranks higher than one with a single issue. Each score includes a trend arrow showing how it has changed since the last scan.




FIGURE 2 How the signal score works

Each score is displayed with one of five color bands that represent score ranges, and flagged items keep their badge regardless of the band they fall into.

Each preset view belongs to a different stakeholder. **Capacity** weights staleness most heavily, helping storage admins find tiering, archive, and reclaim targets. **Risk** prioritizes external access and entitlement concentration, helping security analysts quickly spot potential vulnerabilities. **AI-ready** emphasizes file type density and freshness, which surfaces the best candidates for model training or a deeper classification scan for the data team. **Balanced** provides a starting point for most users.

Switching preset views re-ranks the estate instantly without triggering another scan. Weights are admin-editable, and every export records which preset view produced it. The same folder can score 58 under balanced and 91 under risk.

From your entire estate to a single file

One hierarchy covers every source (location, repository, folder, and file). Each level carries the name your team already uses, so a share stays a share, a bucket stays a bucket, and a SharePoint site stays a site. There’s no need to learn a translation table to navigate your own storage.




FIGURE 3 Overview of an entire estate

The landing view shows a nested tree map. Locations are sized by capacity, with each one showing its largest repositories inside, color-coded by signal score. Every tile shows its name, capacity, file count, native type, score, and trend arrow. Hover over a tile to see the evidence behind the score, including all four signal metrics as bars, the top owner, and any flags. Anything below the top 100 rolls into another tile that opens as a ranked list, ensuring nothing is hidden.

You can find a folder by owner, path, or any other metadata attribute from one search bar with AND, OR, and NOT. Everything on the screen carries the timestamp of the scan it came from, so you always know how current the answer is.

The hotspot for deeper scan

The hotspot for deeper scan gathers the repositories and folders with the highest signal scores, making it easy to focus on the areas most likely to benefit from investigation.

Select the items that matter and add them directly to a scanning policy. That's the whole path from suspicion to scheduled scan, with no ticket or second tool required.

From signal to action

- Export what you find as a .csv or .xlsx file on demand.
- Fire a tiering job or a cleanup ticket without opening the console by feeding the intelligence into your own dashboards, ticketing, or automation over the API or MCP.
- Escalate a source for classification directly from the drill-down, one at a time or in groups built from the scan results.

All of this comes off the same scan. Everpure Data Intelligence adds meaning in layers on top of it by identifying document types, extracting entities, and mapping relationships in a knowledge graph that your teams and AI systems can query directly.

[Learn More About Everpure Data Intelligence](#)

[Visit Our Website](#)

[800.379.PURE](#)

