Diskover

MARKET BRIEF

Buying capacity got hard. Finding it did not.

The next petabyte you need is already in your estate.

Western Digital has sold its entire 2026 output to hyperscale AI buyers, and nearline lead times now run as long as two years. What is missing is not capacity. It is an inventory of the capacity you already have.

100%
of Western Digital's 2026 hard drive output is already committed

2 YRS
the longest current wait for new nearline capacity

36%
of stored data is considered dark or unclassified

88%
of files not accessed in more than a year

At a glance

KEY CHALLENGES

Enterprise SSD contract prices rose about 80% in Q1 2026 alone, and Gartner estimates NAND flash prices will rise 234% this year. A 2026 refresh budgeted at 2024 prices will not buy the same capacity.

THE SOLUTION

Diskover indexes every file across NAS, cloud, object, and archive into one catalog, reading metadata only. It shows what each tier is holding, how old it is, who owns it, and what has gone untouched for years.

WHY IT MATTERS

Capacity you reclaim arrives this quarter. Capacity you order arrives next year. When supply is allocated and priced at a premium, the estate you already own is the only storage with a short lead time, and every terabyte you recover is one you don't have to buy at 2026 prices.

Core capabilities

INDEX EVERY TIER

Distributed crawlers read SMB, NFS, S3, Azure Blob, tape, and archive into a single catalog, scaling out across billions of files. Nothing is copied, and nothing is moved to build it. Your data stays exactly where it is until a workflow says otherwise.

SEE WHAT YOU ARE PAYING FOR

Every file arrives with its owner, size, location, when it was last modified, and when it was last read. Those two dates are rarely the same, and the gap between them is where reclaimable capacity hides. That is the starting set, not the limit. Plugins extend the catalog with the fields your business actually runs on.

FIND THE RECLAIMABLE

Surface duplicate, stale, orphaned, and abandoned files, ranked by the capacity they occupy and the cost of the tier they sit on. The result is not a list of problems. It is a sized opportunity, with the largest and most expensive candidates already at the top.

ASK IN PLAIN LANGUAGE

Ask a question in plain language, and the AI Connector turns it into a query against the catalog's metadata, then returns the matching files. It also proposes what to do next, whether that is to move, tier, or retire. AI-assisted. Human-approved, every time.

ACT ON POLICY

Turn an approved decision into a standing rule. Files move, tier, retain, or clear on the schedule you set, and the same rule runs again next quarter without anyone having to rebuild it. Every action is recorded as it happens, so the catalog keeps describing the estate accurately.

KEEP THE EVIDENCE

Every decision carries its evidence: what the file was, who owned it, when it was last touched, and which rule applied. When an audit, a retention question, or a policy review arrives, the answer is already on record rather than reconstructed from memory.

About Diskover

Diskover helps enterprises structure unstructured data across cloud, on-prem, and edge environments, making it easier to search, govern, optimize, and deliver file data for compliance, analytics, and AI-ready workflows.

Contact

diskoverdata.com
sales@diskoverdata.com
(800) 560-5853
San Jose, CA

How it works

1 CONNECT

Point Diskover at your storage. Distributed indexers read metadata from each system you give them, at whatever scale it runs. Nothing is copied or migrated to get started.

2 CATALOG

Every file lands in one searchable catalog with its owner, size, location, last modified date, and last read date. Plugins add business context from the systems you already run.

3 CURATE

Rank what is duplicate, stale, orphaned, or abandoned. Decide what to keep, what to move to a cheaper tier, and what to retire, with the evidence attached to each decision.

4 ACTIVATE

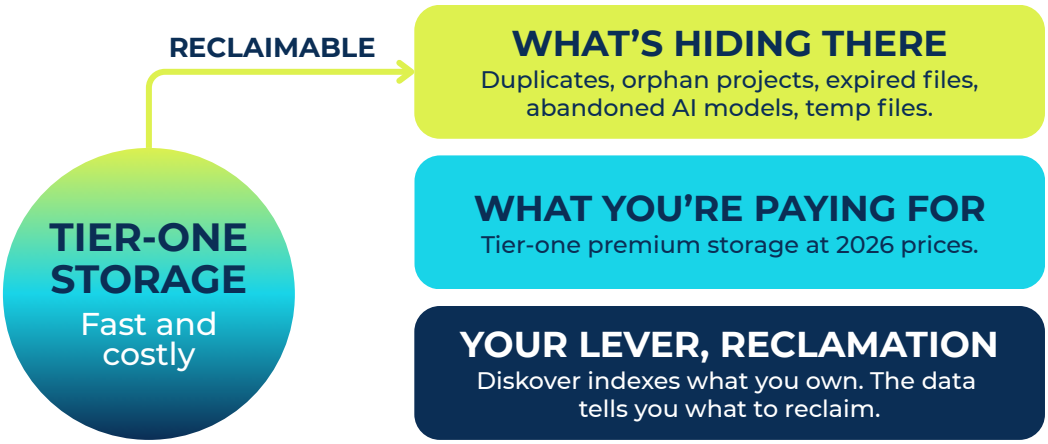
Run the decision as policy. Files move, tier, or clear on schedule, and the catalog records what happened and when.

WORKS ACROSS Cloud · On-prem · Hybrid · NAS · SAN · Object · Tape · Archive · Data lakes · Lakehouses · AI/BI/ML

Built for modern enterprise data environments

No enterprise estate is one system. It is whatever was bought, inherited, or acquired over fifteen years, run by different teams, under different naming conventions, with a decade of project folders in between.

The instinct is to consolidate first and understand it later. Diskover reverses that. It describes the estate exactly as it is, reading only metadata, which is what lets it run against billions of files without moving a byte or asking the storage to do anything it was not built to do.



Common use cases

DEFER THE REFRESH

Size the reclaimable capacity on tier one before you commit to a hardware purchase, so the order you place is the one you actually need.

ANSWER THE RETENTION QUESTION

Produce a file-level record of what is held, where it sits, and how long it has been kept, ready for the moment compliance asks.

CLEAR THE AI PIPELINE

Find the abandoned models, checkpoints, and temporary output that collect around AI projects, and recover the premium storage they have been quietly occupying.

Outcomes teams can feel

CAPACITY THIS QUARTER

Reclaimed space is available the moment the policy runs. Diskover customers reclaim 30 to 50% of their storage by tiering and removing ROT data.

A NUMBER FINANCE CAN USE

Recoverable capacity, priced against what storage costs today rather than when the budget was written. A figure finance can defend, line by line.

FEWER HOURS SPENT LOOKING

Teams ask for what they need and get an answer, instead of walking through folder trees or filing a ticket and waiting on someone else.

GET STARTED

Find out what you can stop paying for.

A 30-day proof of concept, sized to your storage and your prices.

Book a demo

Get the eGuide

Sources: 100%: MLQ.ai, March 2026. 2 yrs: PCSP, September 2026. 36%: Ponemon Institute for Pure Storage, December 2025. 88%: CTERA, September 2026. SSD prices: TrendForce, June 2026. NAND: Gartner, April 2026.

About Diskover

Diskover helps enterprises structure unstructured data across cloud, on-prem, and edge environments, making it easier to search, govern, optimize, and deliver file data for compliance, analytics, and AI-ready workflows.

Contact

diskoverdata.com
sales@diskoverdata.com
(800) 560-5853
San Jose, CA

Revised November 2026