Diskover

SOLUTION BRIEF | ENERGY, OIL AND GAS

From seismic to first oil, every file accounted for.

One enriched metadata catalog for exploration, simulation, and production data across HPC, NAS, object, cloud, and tape.

30%+ of premium HPC storage reclaimed from derivatives and failed runs

1.4B files indexed in two hours, that's ~200K files per second

$1.5M in storage savings at a major global energy company

At a glance

KEY CHALLENGES

Seismic surveys, reservoir simulations, and well programs spawn terabytes of iterations and derivatives at every stage, spread across HPC, NAS, cloud, and tape with no shared catalog. Premium HPC tiers fill with failed runs and temporary files, expired leases and dissolved partnerships demand proof of purge within statutory windows, and geoscience teams wait weeks for clean data before a model can train.

THE SOLUTION

Diskover indexes every file across your entire data estate into one catalog and enriches it with context, such as project, well, and survey, that exploration and production tools leave in file names, extensions, and folder structures. Teams find the right dataset in seconds, tier cold derivatives by policy, purge with an audit trail, and curate seismic and well data for AI.

WHY IT MATTERS

Exploration, processing, and decisions move on schedule instead of through tickets. Premium HPC capacity comes back from derivative and failed-run waste, so the next expansion is smaller or later. Lease and partnership cleanup runs in minutes with proof ready for regulators. Curated seismic and reservoir datasets reach AI and analytics pipelines in days, not weeks.

Why Diskover

ONE CATALOG FROM HPC TO TAPE

Index every file across HPC scratch, NAS, object, cloud, and tape into one searchable view, by fields such as survey, well, format, and age, so one search covers raw seismic, derivatives, simulation runs, logs, and archive, with data staying in place.

DEFENSIBLE DELETION

Tag every file tied to a lease, block, or partnership, run the purge as a policy with a preview, and keep an audit trail of what was removed, when, and by whom, so proof is ready inside the statutory window.

ENERGY FLUENCY IN THE METADATA

Diskover recognizes the file types and project structures that Petrel, Eclipse, Kingdom, and GOCAD leave behind, so SEG-Y and .zgy derivatives, .egrid simulation backups, LAS well logs, and transient .tmp files are identified by metadata such as extension, path, and age.

DIVESTITURE AND M&A READINESS

Locate, validate, tag, and package the datasets attached to an asset or basin under deal deadlines, with an inventory the buyer and the auditor can both read. One energy customer located its divestiture data in under an hour, not weeks.

HPC-AWARE TIERING

When a run completes or fails, policy moves cold derivatives, failed-run outputs, and old iterations off premium HPC storage to NAS, object, or archive and keeps the catalog entry, so scratch capacity comes back without scripts, tickets, or someone remembering.

ASK IN PLAIN LANGUAGE

Ask the catalog "which simulation outputs older than two years are still on HPC scratch" and get matching files back with a recommended action to approve, such as tier, purge, or package. AI-assisted, human-approved.

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

Diskover

SOLUTION BRIEF

How it works

1 CONNECT
STORAGE OR IT TEAM

Point Diskover at HPC file systems, NAS, object, cloud, and tape libraries. Indexing runs in place, so seismic volumes and simulation outputs stay where compute expects them.

Example: a seismic project share is cataloged where it sits.

2 CATALOG
AUTOMATION

Every file becomes a searchable record with rich metadata, such as name, path, and owner, enriched with project and well tokens from the path and classification tags.

Example: a SEG-Y volume now shows its survey, well, owner, and age.

3 CURATE
ANY AUTHORIZED TEAM

Tag by fields such as survey, well, lease, and status. Build saved searches for what teams ask weekly, such as failed runs still on scratch, derivatives past retention, or files tied to a dissolved partnership.

Example: a geoscientist or data manager saves "failed runs" as a search.

4 ACTIVATE
BY POLICY

Run tiering, purge, and packaging by policy with a full audit trail, and deliver curated seismic, log, and simulation metadata to analytics and AI pipelines through the API.

Result: at one energy company, 9 PB was inventoried ahead of a $60B merger.

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

Every survey, well, and run in one view.

Seismic survey
raw SEG-Y, .zgy volumes, processing iterations
HPC

Simulation
Petrel and Eclipse projects, .egrid runs, failed runs
HPC

Production
LAS well logs, field sensor data
NAS

Archive and legal
leases, partnerships, legacy tape
TAPE

Diskover

ENRICHED METADATA CATALOG
Examples: survey · well · block · lease · project · format | from files, paths, plugins, and tags
Every stage of the pipeline and every team, from geoscience to legal, read from the same inventory.

Clean derivatives and failed runs
Found by fields such as extension, path, and age, tiered off HPC by policy.

Purge with proof
Lease and partnership data tagged, removed, and logged for the regulator.

Curate high-value datasets for AI
Clean seismic, log, and simulation datasets reach pipelines in days.

Key outcomes

FASTER TIME TO FIRST OIL

Geoscientists and reservoir engineers find the right survey, run, or log in one search instead of tickets and email, so exploration, processing, and decisions move on schedule.

HPC THAT STAYS CLEAN

Storage teams reclaim 30% or more of premium HPC tiers by moving derivatives and failed runs off scratch by policy, so the next capacity purchase is smaller or later.

COMPLIANCE AND AI-READY

Lease and partnership cleanup runs in minutes with proof of purge ready, and curated seismic and well datasets reach AI and analytics pipelines in days instead of weeks.

GET STARTED

Ready to see it on your estate?

Start with a 30-day proof of concept on your own data, and recover your investment in 3 to 6 months.

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