Diskover

DATASHEET

Diskover + NetApp

One catalog across ONTAP, StorageGRID, and everything around them.

Diskover indexes file and object metadata in place across ONTAP, StorageGRID, and the wider estate, then enriches every record with owner, project, age, usage, and your own attributes. The result is one catalog that answers in seconds. Your files stay on NetApp.

At a glance

WHAT DISKOVER DOES

Indexes metadata from ONTAP and StorageGRID through the protocols they already serve, extends the same index across the wider estate, and enriches each record with the business context you can act on. One schema covers every vendor.

WHAT NETAPP DOES

Unified storage for files and objects. ONTAP serves NFS and SMB at the performance the workload demands. StorageGRID holds S3 object storage at capacity scale across sites. Together, they are the foundation the estate runs on.

BETTER TOGETHER

Capacity planning, retention policy, and AI dataset selection all run on evidence spanning your NetApp systems and the storage beside them, and the recovered capacity funds the next expansion of the NetApp footprint.

WHERE IT FITS

Enterprises running unstructured estates on ONTAP or StorageGRID, with cost pressure, data spread across several platforms, an AI or analytics initiative, a governance mandate, or a refresh on the horizon.

How it works

1 CONNECT
STORAGE OR IT TEAM

Diskover reads ONTAP over NFS and SMB and StorageGRID over S3, with native connections to Azure and Google Cloud.

Example: ONTAP and StorageGRID are cataloged where they sit.

2 CURATE
AUTOMATION

Scheduled scans build one catalog above the storage. Enrichment plugins add attributes in the same pass.

Example: a project file now shows its owner, age, and usage.

3 AUTOMATE
BY POLICY

Policies run on the files a query returns: tagging, tiering, archiving, and retention, with each run logged for audit.

Example: files untouched for three years are archived by policy.

4 ACTIVATE
ANY AUTHORIZED TEAM

Use the AI Connector, the interface, the REST API, or MCP. Curated datasets feed analytics and the AI Data Engine.

Example: an analyst asks which departments drive capacity growth.

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

Business outcomes

VISIBILITY AND COST RECOVERY

Diskover attributes consumption to users, groups, projects, and age bands across the whole estate, identifies redundant, obsolete, and trivial data, and enforces quota and lifecycle policies from a live catalog. Chargeback reports build from the same rows.

GOVERNANCE AND RETENTION

The catalog makes ownership, age, location, and naming patterns searchable in one place, so retention and deletion apply to real files, and audit requests return a defensible set in minutes. Diskover reads metadata and adds file-level actions.

AI AND ANALYTICS READINESS

Diskover identifies, classifies, and curates datasets enriched with business context, so preparation moves from months of manual search and validation to a query against the catalog. Data teams spend their time on model quality instead of data archaeology.

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.

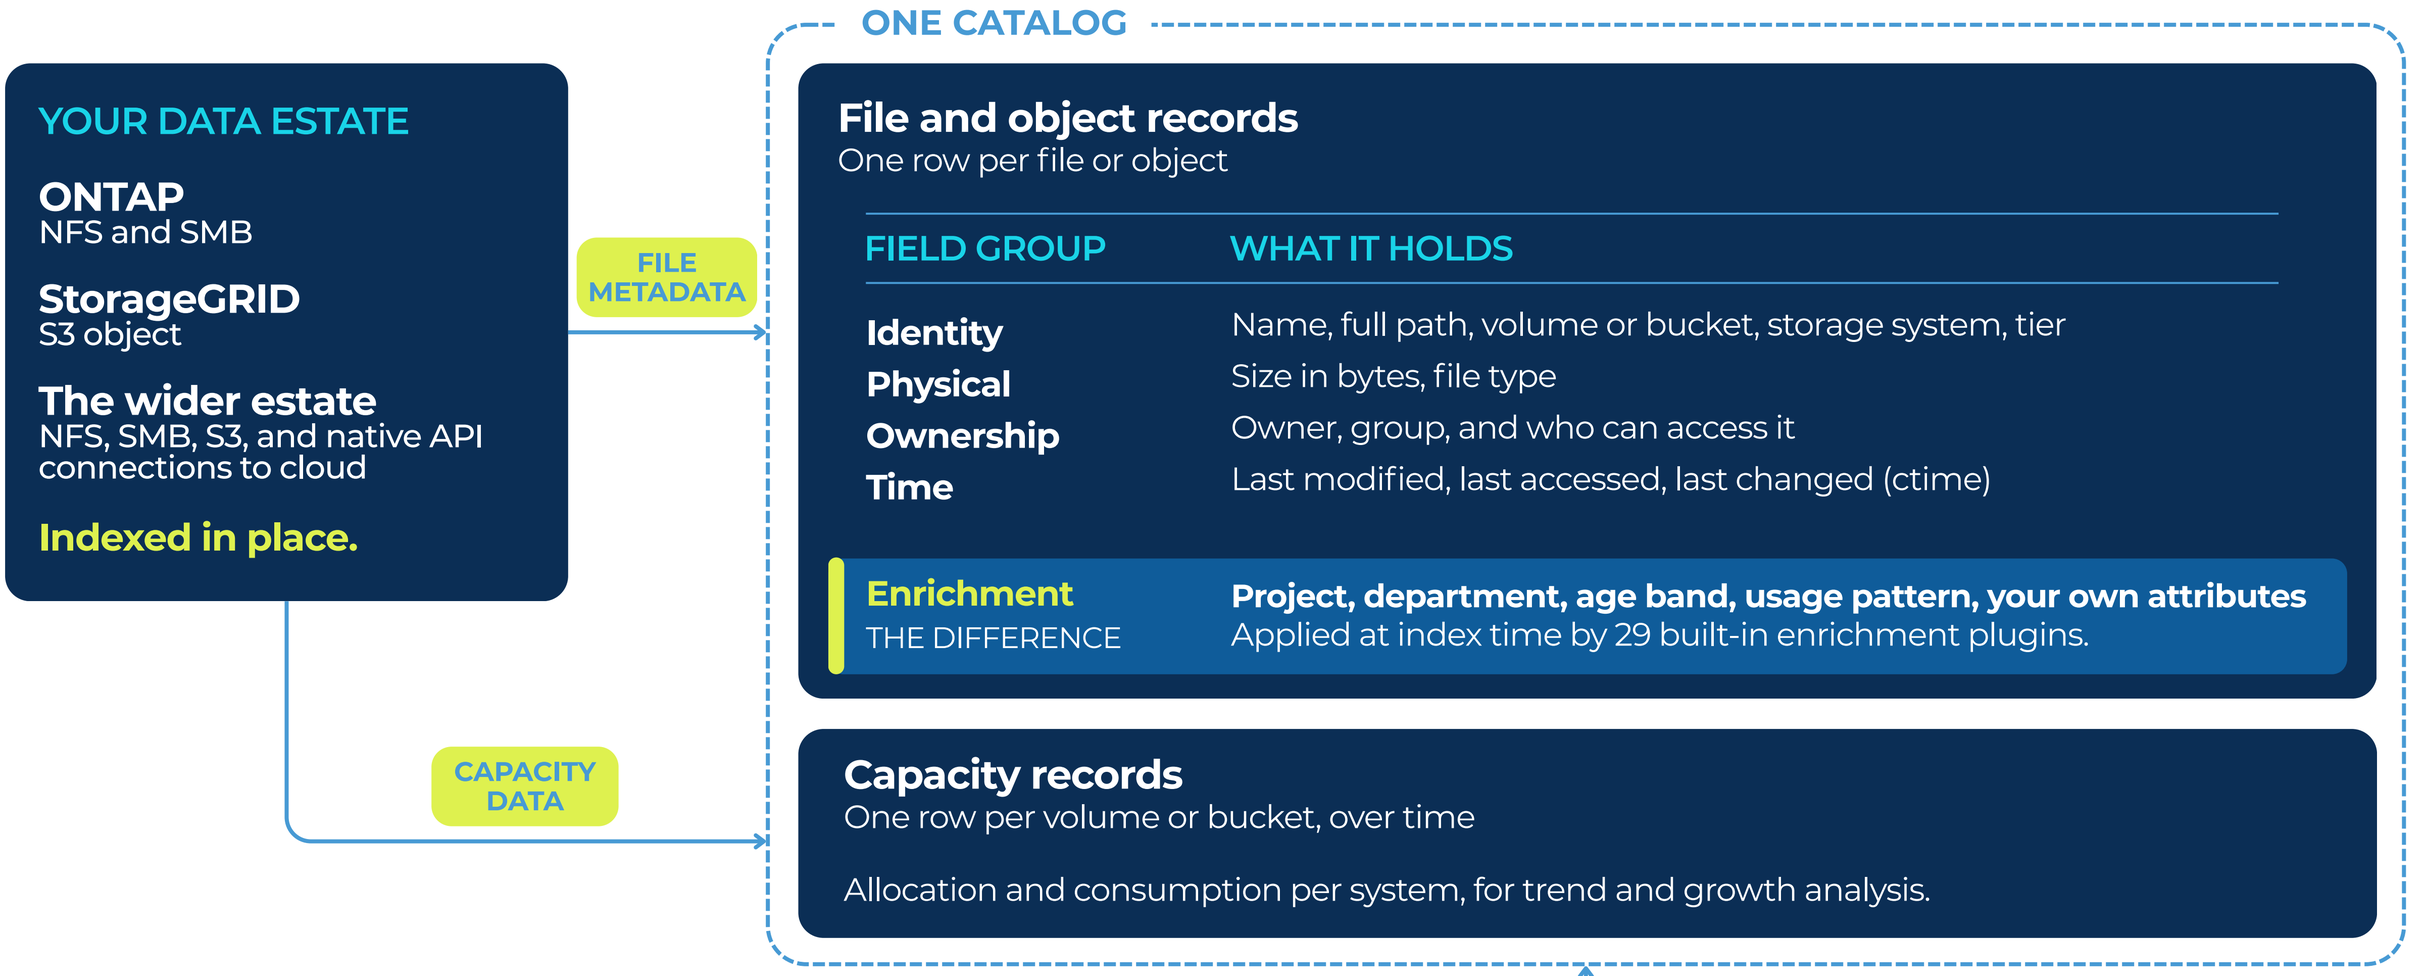
About NetApp

For more than three decades, NetApp has helped the world's leading organizations navigate change – from the rise of enterprise storage to the intelligent era defined by data and AI. Today, NetApp is the Intelligent Data Infrastructure company, helping customers turn data into a catalyst for innovation, resilience, and growth.

Contact

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

What the catalog holds



What you can ask it

Ask in plain language through the AI Connector, or filter in the interface. Answers come from metadata, so they return in seconds, and each one carries a recommended action your team approves before it runs.

- Which volumes and buckets hold data untouched for three years, who owns it, and what reclaiming it returns.
- Which files already qualify for a training run by project, date, size, and format, and how big that dataset is.
- Which files carry naming and location patterns that signal retention or litigation risk, and who is accountable.
- Which departments are driving capacity growth across the whole estate, quarter over quarter, and by how much.

Deployment and support

DEPLOYMENT

Diskover runs entirely inside your environment, on-prem or in your virtual private cloud. A metadata control plane scales horizontally as the estate grows, dispatching task workers onto virtual machines positioned close to the storage each one indexes.

ONBOARDING

Guided onboarding covers package delivery, setup, volume discovery, and dashboards. Insight arrives the same day the storage is connected, and scheduled scans keep the catalog current from there. Adding a volume later takes a connection and a scan.

DATA HANDLING

Diskover reads metadata across every connected system. File content stays on your storage throughout scanning, enrichment, and querying. Movement happens as a deliberate action, executed against the catalog, when a workflow you defined calls for it.

SUPPORT

NetApp provides first and second line support, so you stay inside the support experience your team already runs. Diskover covers third and fourth line escalation. You keep one support relationship rather than adding a second vendor to call.

GET STARTED

Ready to see what your NetApp is holding?

Your whole estate, in one catalog.

Scan to book a demo

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.

About NetApp

For more than three decades, NetApp has helped the world's leading organizations navigate change – from the rise of enterprise storage to the intelligent era defined by data and AI. Today, NetApp is the Intelligent Data Infrastructure company, helping customers turn data into a catalyst for innovation, resilience, and growth.

Contact

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

Revised November 2026