Diskover

CUSTOMER SUCCESS STORY

Entertainment leader reclaims 8 PB, saves $10M.

How a $180B entertainment leader indexed 260 PB across six data centers, found its redundant and outdated media, and cleared 8 PB.



260 PB	8 PB	$10M	$20M+
of media indexed into one catalog	reclaimed across six data centers	saved on storage spend	in CapEx savings

CUSTOMER PROFILE

A global entertainment company valued at more than $180 billion, with studios, broadcast networks, streaming services, and theme parks.

Production, post-production, broadcast, and archive teams create media at petabyte scale across six data centers. A team of 30+ storage, infrastructure, and data platform specialists runs Diskover across that estate.

At a glance

Every production leaves media behind: theatrical assets, broadcast content, sound archives, and IMF deliveries, often in several copies. Across 260 PB and six data centers, the company could see storage filling up, but not which files were still needed. Engineers relied on manual searches and one-off reports.

Diskover indexed more than 100 billion files into one metadata catalog, with the media staying in place. A plugin added production status to each record, and Smart Searches sized redundant, outdated, and trivial content by project, location, age, and size. Tags turned the results into cleanup sets that scheduled workflows archive or remove.

The company reclaimed 8 PB, saved $10M on storage spend, and recorded more than $20M in CapEx savings. The same catalog now feeds Snowflake and Cortex AI.

Facts

INDUSTRY
Media and entertainment: a global entertainment leader.

DATA
Theatrical assets, production media, broadcast content, sound archives, and IMF deliveries.

STORAGE INDEXED
260 PB and 100B+ files on Dell PowerScale, NetApp, and AWS S3.

USERS
30+ storage, infrastructure, and data platform specialists.

LOCATIONS
Studios, networks, and business units worldwide, served from six data centers.

DEPLOYMENT STAGE
Production: cleanup workflows run on a schedule across the estate.

METADATA LAYER
One catalog serves cleanup and feeds Snowflake and Cortex AI.

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

Diskover

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

> *Every storage leader knows the waste and inefficiencies exist. The challenge is quantifying them. Diskover gives them a number they can defend to the CFO – and the insight and confidence to act.*
>
> **Tony Chidiac, CRO**
> **Diskover**

CHALLENGES

SIX SITES, NO SHARED VIEW

Media sat across six data centers and several file systems, with no single inventory of what existed, where, or how old.

OLD WORK ON FAST STORAGE

Completed productions stayed on primary storage because teams could not confirm which assets were final.

COPIES WITH EVERY RELEASE

Duplicates, transcodes, and outdated versions built up with each release, and storage spend grew faster than the library.

HOURS LOST TO SEARCHES

Engineers ran manual searches and one-off reports to answer basic capacity questions for each data center.

SOLUTIONS

ONE CATALOG, 260 PB

Diskover indexed more than 100 billion files across six data centers into one catalog, from metadata only, with media in place.

STATUS ON EVERY FILE

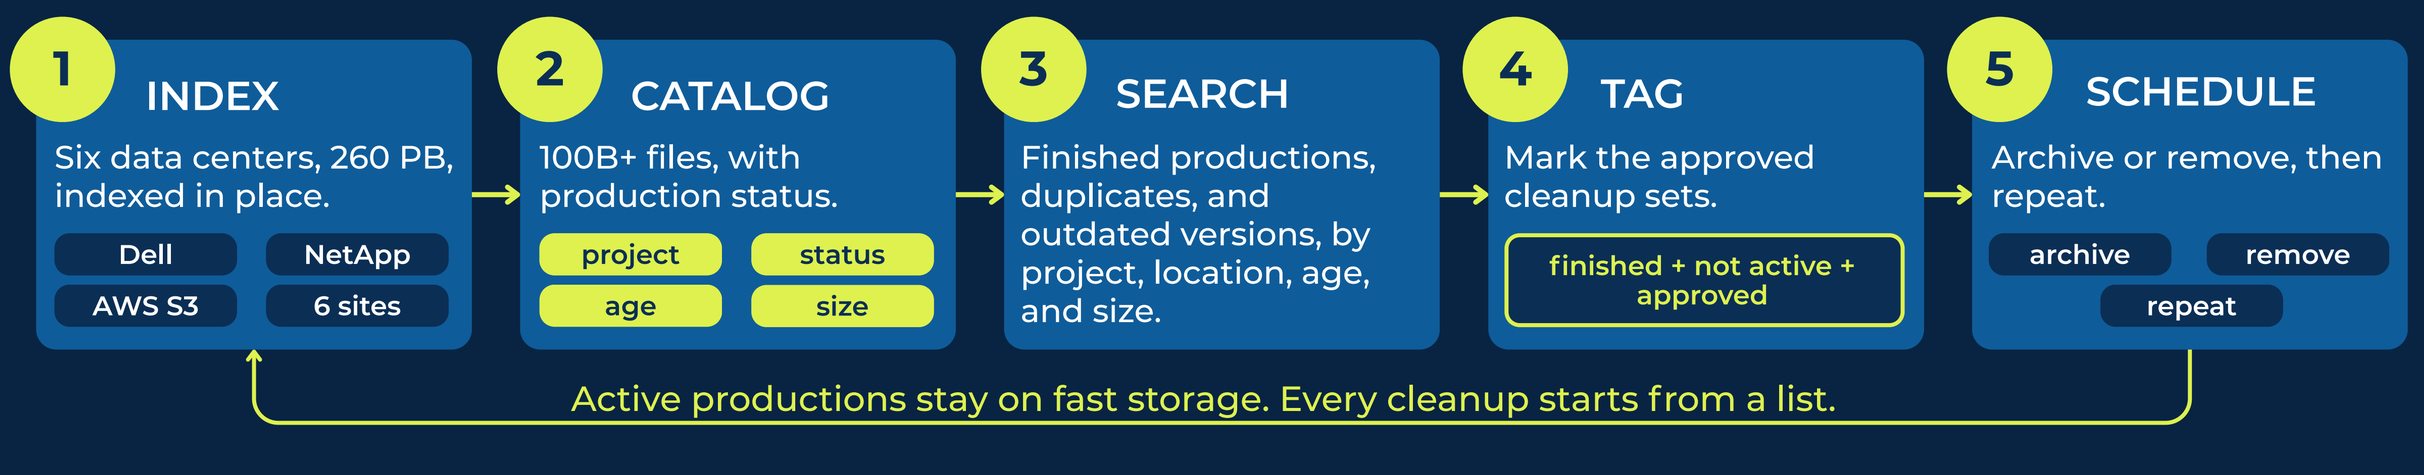
A plugin pulls project and status from the Autodesk production tracking system, so finished work stands apart from active work.

SEARCHES THAT SIZE WASTE

Saved searches on fields such as project, location, age, and size measured the reclaimable content in each data center.

TAG, THEN AUTOMATE

Tags mark the approved cleanup sets, and scheduled workflows archive or remove them, so the estate stays lean.



OUTCOMES

8 PB RECLAIMED

Redundant, outdated, and trivial media cleared across six data centers, returning capacity to active productions.

$10M SAVED

Direct savings on storage spend, alongside more than $20M in CapEx savings.

HOURS BACK TO ENGINEERING

Hundreds of engineering hours recovered each year, as capacity questions became saved searches.

ONE METADATA LAYER

The same catalog feeds Snowflake and Cortex AI and replaced a general-purpose lakehouse workflow.

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.

Diskover

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