Diskover

DATASHEET

Diskover + Snowflake

Metadata first. Content on demand. One catalog across every storage tier.

Diskover indexes file and object metadata in place, then streams a curated subset into your Snowflake account as queryable rows. File content stays on your storage until a query calls for it.

At a glance

WHAT DISKOVER DOES

Indexes every file and object across NAS, SAN, cloud, object, and archive tiers, and describes each record with the business context a query can filter on.

WHAT SNOWFLAKE DOES

Holds those records as governed tables and serves them to Cortex Search, Cortex Agents, CoWork and BI tools under the same Horizon policies as every other table.

BETTER TOGETHER

File records and business records answer the same query. Join what a file is to what the business did with it, inside the account you already govern.

Core capabilities

BUSINESS CONTEXT ON EVERY RECORD

Rules applied at index time attach owner, project, age band, usage, and your own attributes to each record, so a WHERE clause filters on business criteria rather than on path strings.

CONTENT ON DEMAND

Snowflake fetches file content through the Diskover REST API when a query calls for it, from a single document to a bulk load. A Cortex Agent receives the file it requested and only that file.

CURATION BEFORE INGEST

You define what qualifies by file type, age, size, path, and owner. Matching records stream to Snowflake. The rest stay in the Diskover catalog, described and ready if the scope widens.

HORIZON GOVERNANCE FOR FILES

Once file records are Snowflake tables, Horizon Catalog classification, tagging, and access policies apply to them exactly as they apply to your existing tables. One policy engine covers both.

How it works

1 CONNECT
STORAGE OR IT TEAM

Point Diskover at existing storage. Native connections run per storage system, with the access you already have.

Example: NAS, object, and archive tiers are cataloged where they sit.

2 CATALOG
AUTOMATION

Scheduled scans catalog metadata in place across every tier. Plugins add attributes in the same pass.

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

3 STREAM
BY POLICY

Curated records stream into your Snowflake account, optionally through Kafka, and land as queryable rows.

Example: only files from the last 18 months land in Snowflake.

4 ACTIVATE
ANY AUTHORIZED TEAM

Query file records beside business tables, point Cortex Search and Cortex Agents at them, or analyze storage cost.

Example: finance queries capacity growth by department.

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

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.

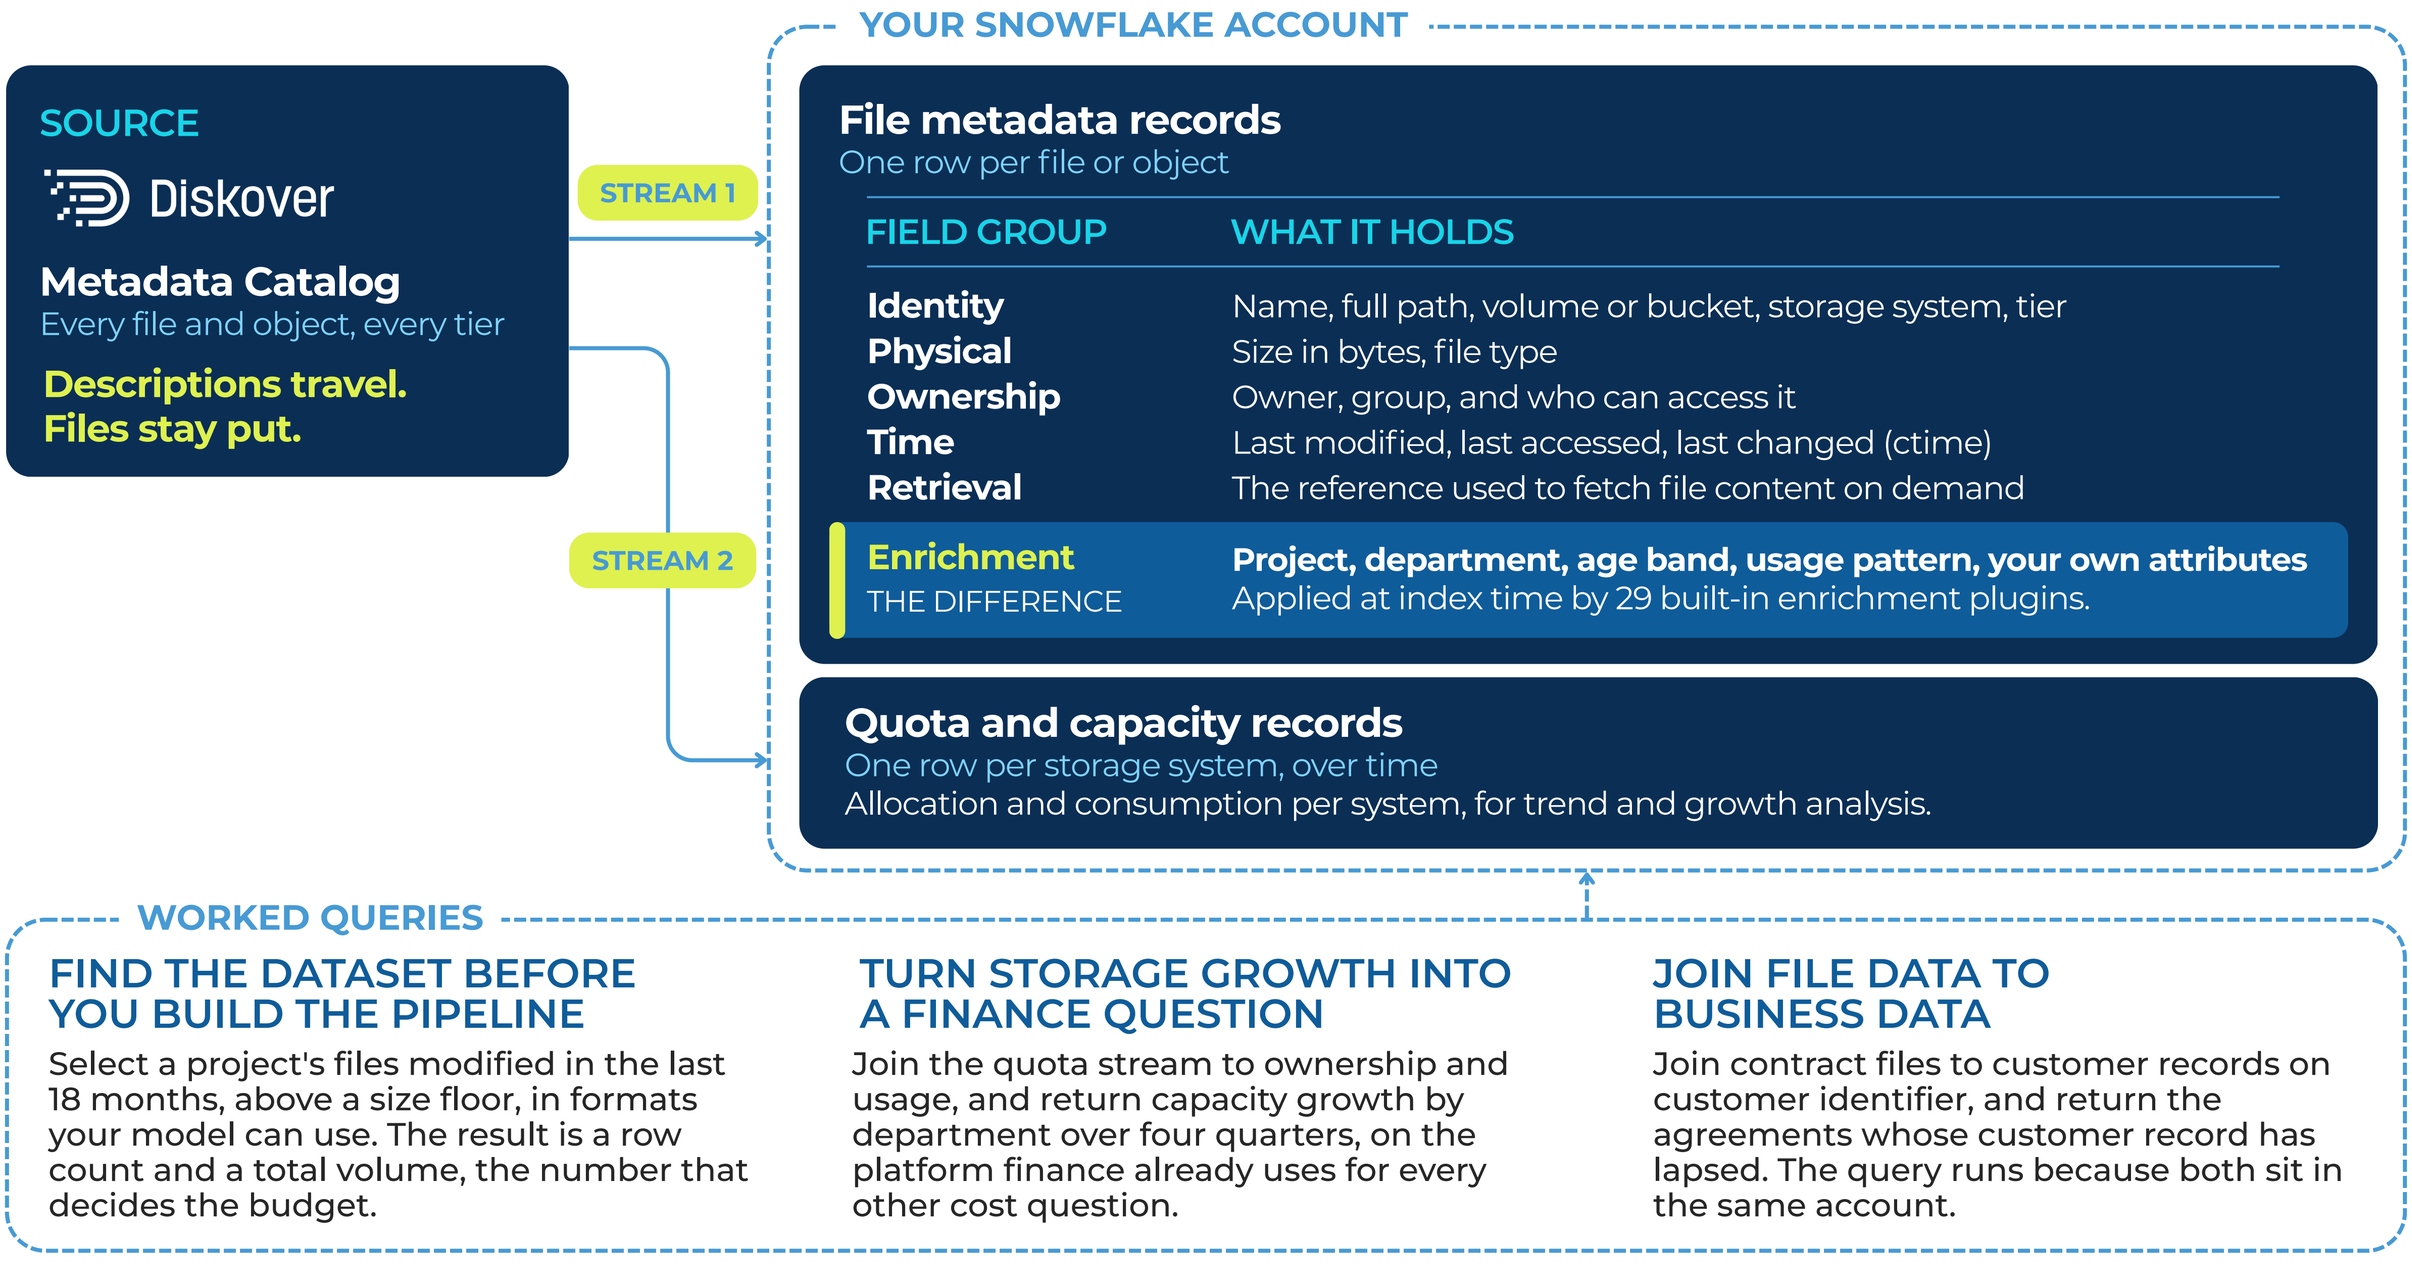
About Snowflake

Snowflake is the platform for the AI era, making it easy for enterprises to innovate faster and get more value from data. More than 13,900 customers around the globe, including hundreds of the world's largest companies, use Snowflake's AI Data Cloud to build, use and share data, applications and AI. With Snowflake, data and AI are transformative for everyone.

Contact

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

Diskover



Business outcomes

PIPELINE COUNT GOES TO ONE

One catalog replaces a connector per storage system. Engineering time moves from maintaining ingestion to answering questions, and a new storage system is a connection rather than a project.

INGEST VOLUME BECOMES A DECISION

The curation predicate runs before transport, so the rows that reach Snowflake are the rows you selected. Ingest volume follows your scope rather than your estate size.

FILE HISTORY BECOMES QUERYABLE

Decades of file history sit beside current business tables, so failure investigations, audits, and asset searches run as queries under the same access controls as the tables beside them.

GOVERNANCE COVERS THE OTHER 80 PERCENT

The file share of enterprise data enters the same catalog, classification scheme, and access policies as your tables, which is the condition for letting an agent near it.

Deployment and security

DEPLOYMENT

SaaS or on-prem, running against the storage you already own at exabyte scale. The request broker is managed for you, and the agents run where the storage is.

CONNECTIVITY

Native connections to NAS, SAN, object, cloud, tape, and archive systems, plus REST APIs and webhooks that keep metadata portable.

INGESTION PATH

Snowpipe Streaming, orchestrated by the shipped connector. Add Kafka for durable transport and replay at scale, and the connector delivers each batch into your tables.

SECURITY

Transfers are encrypted, authenticated, and checksum-verified. Content requests are brokered with scoped, short-lived credentials. Horizon policies govern the rows once they land.

PREREQUISITES

A Snowflake account, agent access to the connector, and connector credentials to write to your tables. A Kafka cluster is optional, for durable transport and replay at scale.

GET STARTED ON SNOWFLAKE MARKETPLACE

Ready to send Snowflake the data it needs?

See what is in the 80% you cannot query today.

Scan to book a demo

Go to Snowflake Marketplace

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 Snowflake

Snowflake is the platform for the AI era, making it easy for enterprises to innovate faster and get more value from data. More than 13,900 customers around the globe, including hundreds of the world's largest companies, use Snowflake's AI Data Cloud to build, use and share data, applications and AI. With Snowflake, data and AI are transformative for everyone.

Contact

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

Revised November 2026