

# MAXIMIZE THE UTILITY OF YOUR DATA LAKE

## Combine the performance and scalability of Snowflake with your data lake to discover all insights



### SUPPORT ALL YOUR DATA

Natively ingest all of your structured and semi-structured data (JSON, Avro, and XML) into Snowflake, or leave that data in your data lake and query it with Snowflake's robust ANSI SQL.



### USE SCHEMA-ON-READ

Use Snowflake's schema-less data ingestion to avoid unnecessary delays when loading data from your data lake to Snowflake. Or, use Snowflake's external tables to access files directly from your data lake.



### CAPTURE DATA UPDATES, FASTER

Configure Snowflake external tables to automatically refresh as the data in your data lake changes. New data is readily available for analysis without expensive and time-consuming overhead.



### ACCELERATE DATA LAKE ANALYTICS

Independently scale storage and compute resources automatically or on the fly. No need to configure or size for peak demand months in advance.



### CAPITALIZE ON NEAR-ZERO OVERHEAD

With Snowflake, avoid the complex architectural changes of configuring, refreshing, and analyzing data within your data lake.



### SUPPORT ALL YOUR USERS

Take advantage of Snowflake's instant and infinite scaling of independent workloads for any level of concurrency without impacting performance.

Cloud object storage has become the de facto choice for data lakes. Even though data lakes offer the ability to store large amounts of structured and unstructured data, gaining all the insights from all that data can be challenging.

When used in conjunction with a cloud-built data warehouse, a data lake can become a treasure trove of insight with very little overhead. Using a standard SQL interface means you can more efficiently discover value hidden within the data lake, and quickly deliver data-driven insights to all your business users.

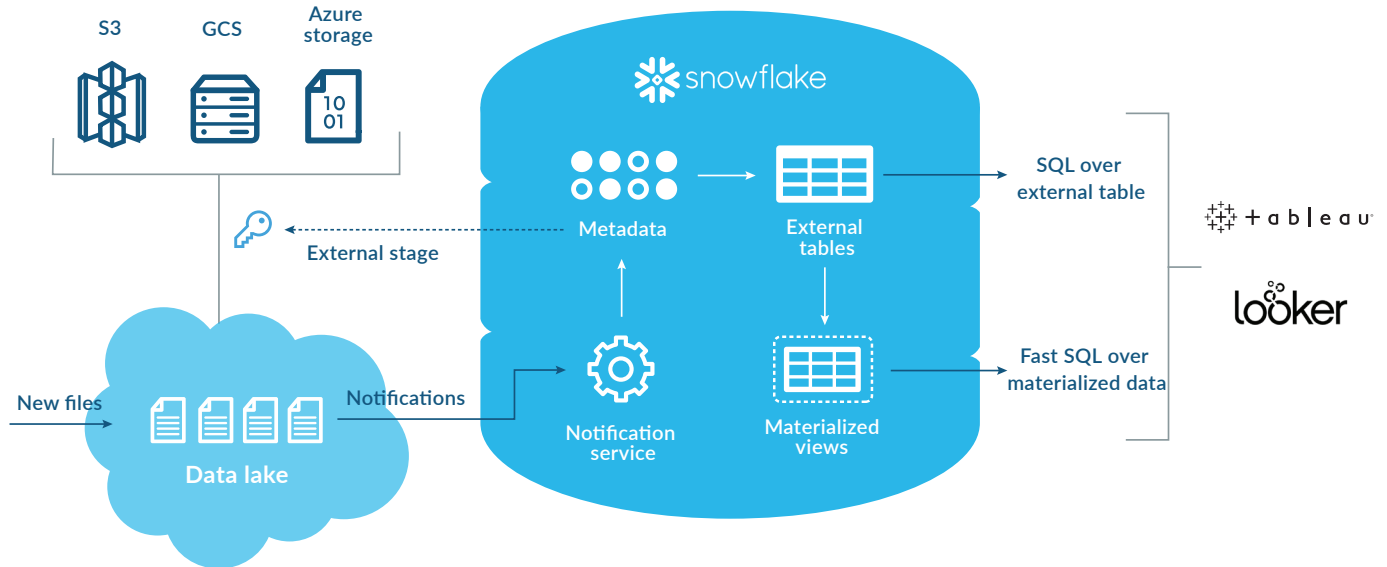
Snowflake achieves this with external tables, which are links to files in a data lake that make the data in it immediately accessible within Snowflake but without the need to lift, shift, or copy data. Materialized views on those external tables preprocess data and therefore accelerates your analytics.

Delivered as an easy-to-use data warehouse service, Snowflake enables you to process and analyze all your diverse data, build multiple databases, query with a common robust ANSI SQL environment, and execute ACID transactional capabilities.

Using Snowflake with a data lake allows you to work with more diverse data sets without requiring you to manipulate data for specific use cases. All the while, you pay for only what you need, when you need it.

Great hype erupted when the data lake first emerged. But that hype soon subsided when the insights from all that data were near impossible to access. Snowflake now delivers on that hype.

## External Stages and Tables in Snowflake



How Snowflake integrates structured and semi-structured data from a data lake

### CUSTOMER SUCCESS



Risk Management Solutions (RMS), the world's leading catastrophe risk modeling company, struggled with burst compute on its data lake to provide ad-hoc query capabilities for its analysts. The challenge: Due to the variable and unknown schema of their parquet files and query patterns, ingesting terabytes of data into a data warehouse was impractical. With Snowflake external tables, RMS was able to:

- Keep its data in the data lake, treating it as a single source of truth and writing queries on parquet files directly
- Take advantage of Snowflake's data warehouse, instant elasticity, and high concurrency, while querying external files
- Apply schema-on-read using variant columns in external tables
- Speed up queries using materialized views

### ABOUT SNOWFLAKE

Thousands of customers deploy the Snowflake Cloud Data Platform to derive all the insights from all their data by all their business users. Snowflake equips organizations with a single, integrated platform that offers the only data warehouse built for any cloud; instant, secure, and governed access to their entire network of data; and a core architecture to enable many other types of data workloads, such as developing modern data applications. Find out more at [snowflake.com](https://snowflake.com).