ONEHOUSE
Your Universal Data Lakehouse

Onehouse Iceberg for Snowflake

Truly open data, at 60% lower cost

Apache Iceberg™ is an open table format that is often used to store data in a vendor-neutral format, without lock-in. Snowflake and other commercial data platforms now provide official support for Iceberg, making Iceberg an increasingly common option for users who wish to keep control of their data and avoid proprietary formats.

However, Iceberg still needs to be managed. Tables need to be optimized (clustering, compaction, data cleaning), ETL pipelines must be processed to transform data, tables need to be synced with catalogs, and query engines need access to the tables. All of this can be achieved manually, but most users prefer a fully-managed service to reduce complexity.

Today, Snowflake users have two primary options to deploy a fully-managed service for Iceberg:

Snowflake-Managed (Native) Iceberg Tables

- Fully-managed and optimized by Snowflake.
- ETL pipelines are executed by Snowflake SQL.
- Access requires using the Snowflake Catalog and SDK.

Onehouse-Managed Iceberg Tables

- Fully-managed and optimized by Onehouse.
- ETL pipelines are executed by the Onehouse Quanton engine running on the Onehouse Compute Runtime.
- Data is synced to any of the leading catalogs by OneSync™ multi-catalog sync.

	Table Formats	Table Optimizations	ETL	Catalog	Query Engine
Snowflake-managed Iceberg	Iceberg only	Snowflake	Snowflake SQL	Snowflake Catalog	Snowflake SQL (via Snowflake Open Catalog)
Onehouse Iceberg	Apache XTable to support Iceberg, Hudi, and Delta Lake	Table Optimizer (compaction, clustering, data cleaning)	Quanton + Onehouse Compute Runtime **60% better price/performance**	Multi-Catalog Sync	Supports all popular query engines and warehouses

The Onehouse Advantage: Open Data at 60% Lower Cost

Onehouse delivers the **Universal Data Lakehouse**, the industry's most open and cost-efficient data lakehouse. Onehouse is a fully-managed cloud service that manages your Iceberg tables, provides a serverless SQL environment to run your ETL pipelines, syncs your tables to all the leading catalogs, and connects to all the popular query engines.

Onehouse provides the following advantages compared to Snowflake-managed Iceberg tables:

Truly Open Data

- **No table format lock-in.** Present your tables in Apache Iceberg™, Apache Hudi™, and Delta Lake format with Apache XTable™ (Incubating) providing metadata translation between formats.
- **No catalog lock-in.** OneSync™ auto-syncs metadata to AWS Glue, Hive, Snowflake, Unity Catalog, Google Data Catalog, and others.
- **No query engine lock-in.** Query with Snowflake, Spark, Trino, Databricks, Ray, and any of the popular query engines.

60% lower cost

- More than 50% of cloud data warehouse spend is on ETL pipelines. **Onehouse reduces the cost of running these pipelines by 60%.**
- SQL pipelines are processed by the Quanton execution engine running on the Onehouse Compute Runtime (OCR). Quanton and OCR are **optimized for ETL workloads and provide a 60% price/performance advantage** compared to running SQL pipelines on cloud data warehouses.

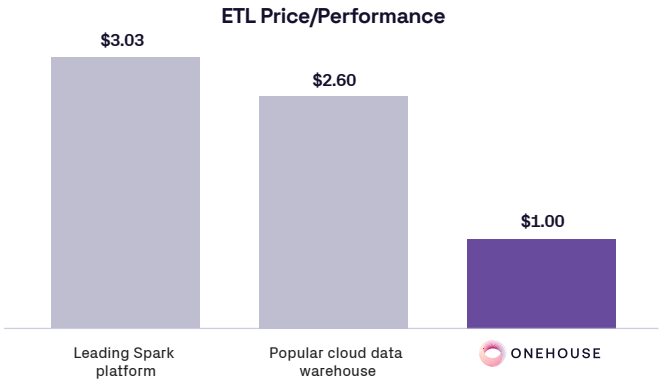

Performance parity with Snowflake

- Ingest with minute-level data freshness from CDC and streaming data sources.
- Keep tables at peak performance with Table Optimizer, which automates compaction, clustering, and data cleaning.
- Query your tables 2-10x faster than with unmanaged Iceberg, and with similar performance to Snowflake-managed Iceberg.

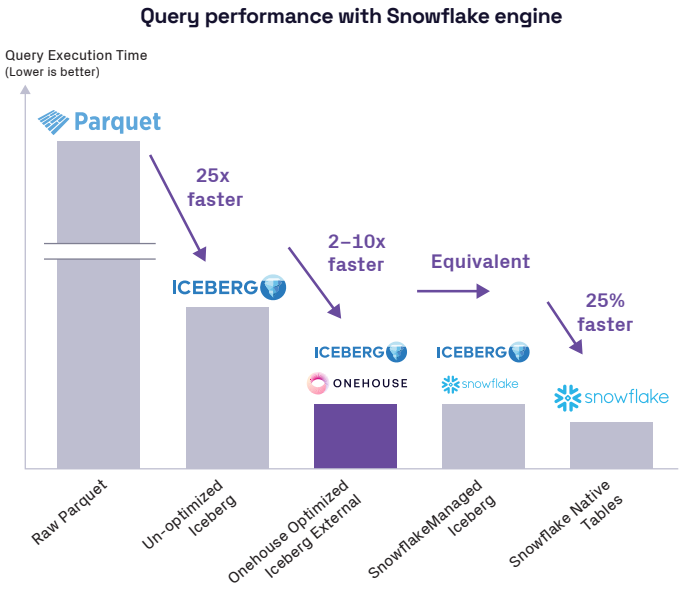

Find out more at onehouse.ai/snowflake

ONEHOUSE