



ONEHOUSE
Your Universal Data Lakehouse

Lightning-fast ingestion, incremental transformations, and intelligent optimizations.

The only open data platform for all workloads from analytics to GenAI.

The reality of most data platforms



Slow

50% of engineering time is spent on maintenance instead of delivering insights.



Expensive

30% of cloud spend is wasted on unused resources, duplication, and inefficient compute.



Locked-In

Vendor-controlled storage & compute limit flexibility, while 65% of enterprises still manage hybrid or on-prem infrastructure.

Onehouse is the only fully managed cloud data lakehouse that delivers



2x Faster

Lightning-fast ingestion, ETL, and queries.



1/2 the Cost

Slash compute and storage expenses.



1/5 the Engineering Effort

Automate optimizations and reduce maintenance.



Unlimited Interoperability

Work across table formats, catalogs, and engines.

From the original creator of Apache Hudi™ adopted industry-wide by the largest data lakes

Powered by open source technologies you love

Uber amazon Robinhood PHILIPS

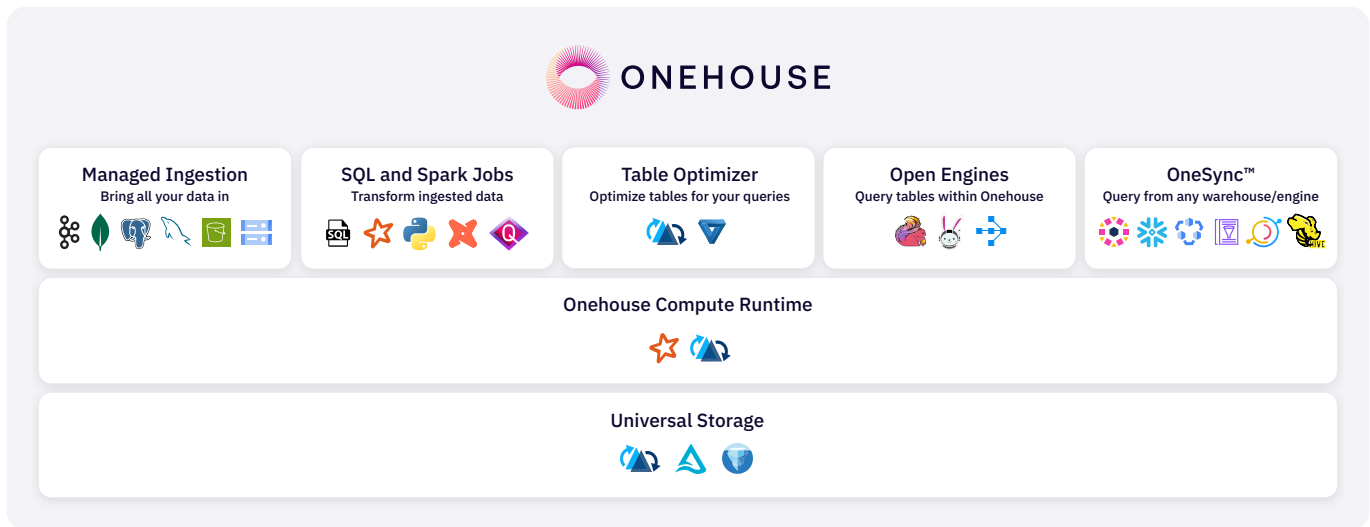
nerdwallet ûdemy zendesk Disney+ hotstar

Tencent Cloud Google Cloud



Built on open source, optimized for enterprise

Onehouse unifies ingestion, optimization, and querying across all major table formats while keeping data independent, performant, and cost-efficient.



- ✓ Fully managed IaaS delivered in days with minute-fresh data
- ✓ Compute runtime optimized for cost-efficient Spark + data lakehouse workloads
- ✓ Universal interoperability across catalogs and open table formats

Quanton Engine for SQL, Onehouse Compute Runtime, and managed clusters have all been hugely beneficial for Conductor in abstracting complexities around cluster management, compute allocation, downtime, debugging, and more. We have integrated our internal frameworks against the Onehouse Spark SQL endpoint and run workloads on terabytes of data to automate data pipeline workflows and data quality checks.

— Uday Rajanna, Principal Engineer, Conductor



From BI to AI with Onehouse's Open Data Lakehouse

Onehouse powers BI, machine learning, and GenAI workloads on a fully open, high-performance lakehouse, eliminating lock-in while reducing costs.



BI / Analytics

Run SQL analytics and BI reporting with Snowflake, Redshift, BigQuery, and Trino directly on your lakehouse.



Machine Learning / Data Science

Accelerate machine learning pipelines with structured, high-performance storage optimized for AI training and inference.



Gen AI / LLMs

Store and retrieve vector embeddings natively for GenAI applications without the need for additional databases.

Find out more at onehouse.ai/onehouse-cloud

