




# How ATRenew Achieves Low Latency Data Lake Analytics With 10x Better Cost-Effectiveness



**10x+**

better cost-effectiveness

**50%**

less compute resources

**5x-10x**

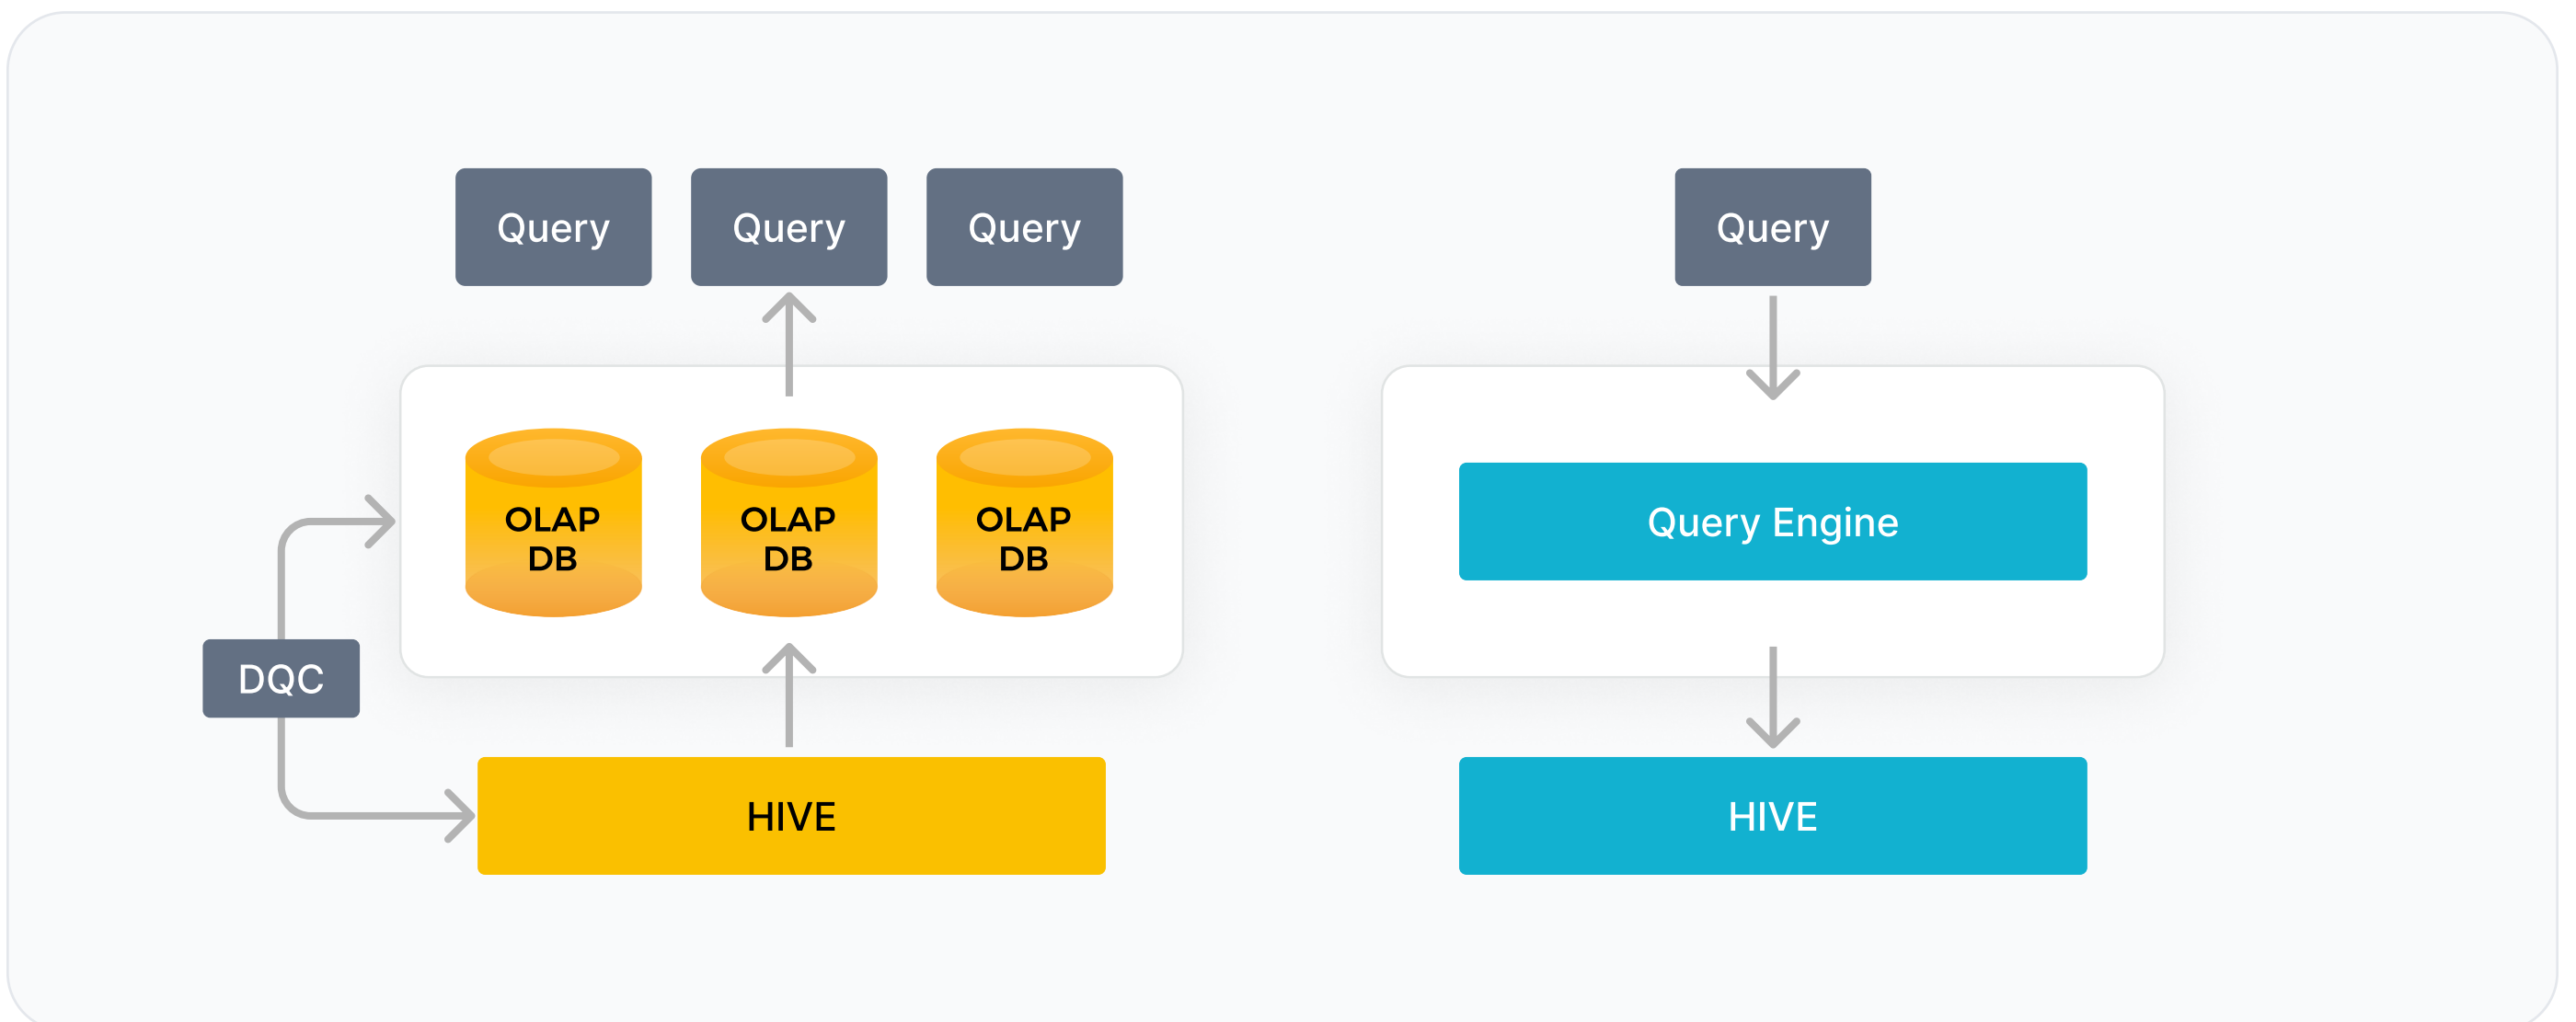
performance gain on 80% of the queries

## About ATRenew's Watcher Platform

ATRenew (NYSE: RERE), established in 2011, has become a frontrunner in the "Internet + environmental protection" sector, promoting a circular economy. Its reporting platform, Watcher, transitioned from Trino to StarRocks as its data lake query engine, enabling low latency queries under hundreds of QPS, resulting in a more than tenfold improvement in cost-effectiveness.

## Challenges

ATRenew handles a large volume of data through its detailed processes of inspecting, grading, and reselling recycled items. Daily, the company gathers data ranging from tens to hundreds of terabytes and maintains a data lake that stores five years of historical data. The sheer volume presents a significant challenge to their reporting platform Watcher due to the profile of queries, which is complex and often requires accessing up to half of the historical data stored.



The financial impracticality of duplicating and ingesting this data into a data warehouse arises from the costs of storing an additional copy of data, the hardware resources needed for rewriting the data into another format, and maintaining the data ingestion pipeline. Hence, directly querying the data lake remained the only viable option for ATRenew.

Initially, ATRenew implemented Trino as their query engine, but it fell short on performance, particularly under Watcher's high concurrency requirements of over 200 QPS, resulting in poor user experience. This is unacceptable for the Watcher platform, the engineers are searching for a query engine that is more optimized for complex queries.

## Solution

StarRocks emerged as a replacement for Trino, its C++ SIMD-optimized execution engine promises superior performance against complex multi-table OLAP queries, which are exactly the queries Watcher was struggling with.

Benchmark tests on TPC-DS 500GB show StarRocks being up to 4.16x faster than Trino with the same hardware resource and up to 1.59 times faster on 1/3 of the hardware resource.

Concurrency	Trino403(s)(9 worker)	StarRocks 2.4.1(s) (9 BE)	StarRocks 2.4.1(s) (3 BE)	Trino/StarRocks 9BE	Trino/StarRocks 3BE
10	3832.00	985.33	2820.00	3.89	1.36
20	8083.67	1952.33	5438.67	4.14	1.49
30	11990.33	2879.67	7554.33	4.16	1.59

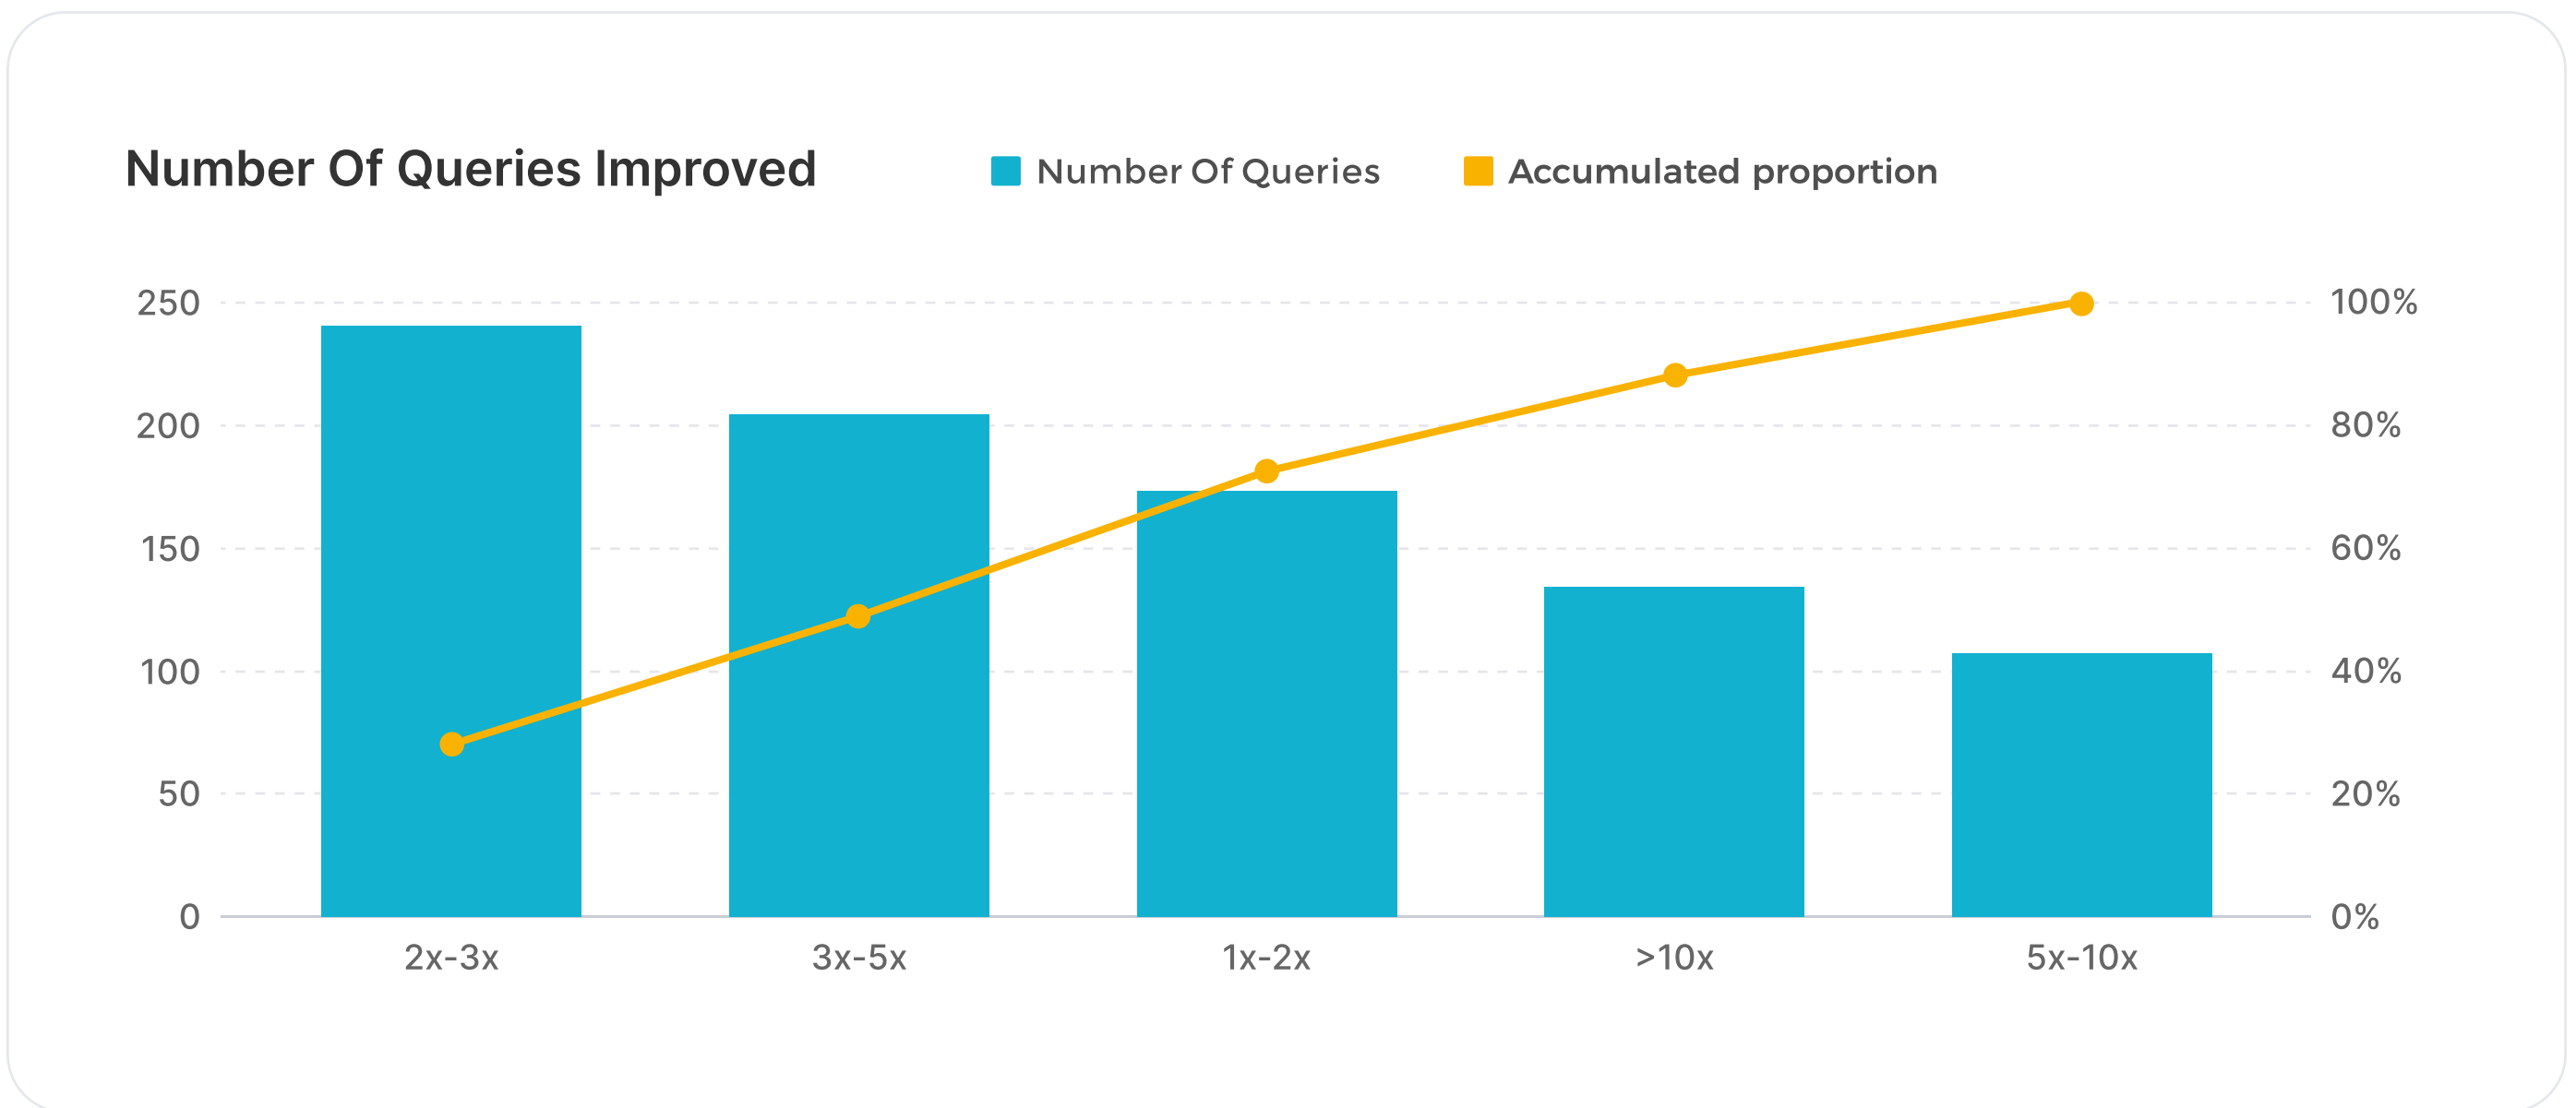
Tests on actual production workloads see greater performance improvements, with up to 16.03x under 20 concurrency.

Concurrency	Trino(s)	StarRocks(s)	Trino/StarRocks
1	1105.33	163.33	6.77
10	2210.00	201.67	10.96
20	2746.33	171.33	16.03

## Result

StarRocks version 3.1 has been successfully deployed in production on ATRenew's Watcher reporting platform, replacing Trino with StarRocks and utilizing only half the number of nodes that were previously used by Trino. StarRocks' full Trino dialect support enabled all of the queries to successfully migrate to StarRocks with ease.





In production, 94% of Watcher's queries saw performance improvements. Around 80% of these queries saw performance increases ranging from 5x-10x. This enhancement not only cut the infrastructure cost by half, it also accelerated decision-making processes and improved efficiency throughout ATRenew.

## What's Next For ATRenew

In light of StarRocks' outstanding performance in production, ATRenew is looking to further explore StarRocks' capabilities and expand its usage in other scenarios:

- Explore StarRocks' shared data deployment to dynamically scale for business fluctuations.
- Explore StarRocks' data cache feature to further accelerate query performance for scan-heavy queries.
- Utilize StarRocks to support business logic ETL (extract, transform, load) jobs on Hive.



[Join the StarRocks Slack Community](#) 

Share notes, ask questions, and get feedback from thousands of your peers working at world-class companies.