

ClickHouse vs. StarRocks

	ClickHouse	StarRocks
Architecture	Scatter-Gather	MPP
SQL Syntax	Limited support	Full support (SQL-99)
Data Ingestion	Can't guarantee once and only once ingestion	Ingest once and only once with ACID on ingestion
Distributed Join	Not supported	Supported
Real-Time Updates	Asynchronous updates, not real time	Synchronized updates and supports UPDATE/DELETE
Concurrency	Limited number of concurrent users	High concurrency with 10,000+ queries per second
Count Distinct	Performance degrades as data volume grows	Supports accurate and approximate count with consistent performance
Query Planning	No cost based optimizer	Built-in cost based optimizer
Scalability	Requires manual data rebalancing when scaling	Data rebalancing is automatic when scaling
Data Partitions	Data partitioned into parts and parts are tied to nodes	Partition + bucketing and not tied to physical servers
Data Lake Queries	Rudimentary support	Supports queries on Hive, Hudi, Iceberg, and Delta
Federated Queries	Not supported	Supports queries on Hive, MySQL, ES, and JDBC sources
Operations	Cumbersome and reliant on Apache Zookeeper	Streamlined architecture with no 3 rd party dependencies

Materialized Views	Asynchronously updated with tables and doesn't support query rewrites or multi-table	Synchronized with tables, supports query rewrites, and materialized views on multi-table
Materialized Views: Query Rewrites	Materialized view requires SQL changes	Materialized view seamlessly supports query rewrites
Materialized Views: Incremental Updates	No support for incremental updates	Supports incremental updates for real-time scenarios

About CelerData

CelerData enables enterprises to quickly and easily grow their business with a unified analytical engine that is 3X the performance/cost of any other solutions on the market. CelerData is the only platform uniquely designed for the next generation real-time Enterprise, unleashing the power of business intelligence to help accelerate Enterprise digital transformation. Used worldwide by market leading brands including Airbnb, Lenovo and Trip.com, CelerData generates critical new insights for these data-driven companies. To learn more, please visit, www.celerdata.com