



July 13, 2023

Presenter:

Albert Wong

Senior Developer Advocate @CelerData

Host:

Shanna Chan

Principal Solution Architect @Red Hat

Stay Connected at the San Diego [Meetup.com](https://www.meetup.com/Red-Hat-San-Diego) group link: <https://www.meetup.com/Red-Hat-San-Diego>

Solving customer problems through partnerships

Today's IT requires an ecosystem of partners that can work together to solve for various use cases.




Red Hat



Red Hat OpenShift and Red Hat Application Foundations

A strategic approach to application development and delivery





Harness the Power of Modern Data Lakehouse

Albert Wong
albert.wong@celerdata.com
Senior Developer Advocate @ CelerData

Agenda

- What is online analytical processing (OLAP)? Why is OLAP important?
- History of OLAP
 - Batch OLAP
 - Data Lakes
- Future of OLAP
 - Data Lakehouse
- Accurate, Real-Time Analytics has always been the goal
 - How StarRocks makes Real-time Analytics Easy

What is online analytical processing?

Online analytical processing (OLAP) is software technology you can use to analyze business data from different points of view. Organizations collect and store data from multiple data sources, such as websites, applications, smart meters, and internal systems. OLAP combines and groups this data into categories to provide actionable insights for strategic planning. For example, a retailer stores data about all the products it sells, such as color, size, cost, and location. The retailer also collects customer purchase data, such as the name of the items ordered and total sales value, in a different system. OLAP combines the datasets to answer questions such as which color products are more popular or how product placement impacts sales.

Why is OLAP important?

Faster decision making

- Businesses use OLAP to make quick and accurate decisions to remain competitive in a fast-paced economy. Performing analytical queries on multiple relational databases is time consuming because the computer system searches through multiple data tables. On the other hand, OLAP systems precalculate and integrate data so business analysts can generate reports faster when needed.


Non-technical user support

- OLAP systems make complex data analysis easier for non-technical business users. Business users can create complex analytical calculations and generate reports instead of learning how to operate databases.

Integrated data view


- OLAP provides a unified platform for marketing, finance, production, and other business units. Managers and decision makers can see the bigger picture and effectively solve problems. They can perform what-if analysis, which shows the impact of decisions taken by one department on other areas of the business.

State of the online analytical processing (OLAP) today



Data Lakehouse

Open data management architecture that combines the flexibility, cost-efficiency, and scale of data lakes with the data management and ACID transactions of data warehouses



Data Lakehouse – Storage Layer

This is the base layer where raw data is stored. The data can be structured, semi-structured, or unstructured.

It's essentially where the "data lake" part of the data lakehouse resides.

It is:

- Extremely scalable
- Cheap
- Durable



Data Lakehouse – Governance & metadata Layer

Data management tasks like data cataloging, data lineage, data access controls, quality checks, and metadata management.

Metadata gives "structure" to your data.

- ACID control
- Indexes
- Versioning
- ...
-

	Iceberg	Delta Lake	Hudi
OSS	Spark, Hive, Flink, Presto	Spark, Hive, Presto	Spark, Flink, Kafka
ACID	Yes	Yes	Yes
Data Format	Parquet, Avro, ORC, CSV	Parquet	Parquet
Current Pointer	Metastore, File system with version File	Transaction log	Timeline commit
Conflict Resolution	Optimistic	Optimistic	Optimistic
Time Travel	Yes	Yes	Yes
Schema Evolution	Yes	Limited	Partial
Partition Evolution	Yes	No	No
Merge	Yes	Yes	Yes
Storage Abstraction	Yes	Yes	No
Language	Java & Python	Scala, Java & Python	Java & Python
Compaction	API based	Manual	Automated

Metastores



Open Table Formats



HIVE

ICEBERG



Apache



Data Lakehouse – Query Layer

Enables users to extract insights from the data stored in the data lakehouse, typically with SQL.

Heterogeneous workload, batch processing and sub-second interactive querying.



Photon

The next generation engine for the Lakehouse

[Get started](#) [Learn more](#)



Future is Open: Open Lakehouse vs Proprietary Lakehouse

Proprietary



Open Storage



Open



StarRocks

an open source data lakehouse

StarRocks: High-Performance Analytical Database



Fully Featured OLAP DB
High Performance
Scalable



Uses the MySQL Protocol




Open Source

CelerData - StarRocks for the Enterprises

Best way to
run
Starrocks

Automated.
Supported.
Secured.



CelerData Key Facts

Headquarters: Menlo Park, CA

Number of Employees: 90

Total Funding: \$60MM

Open-Source Contributors: 200+

Enterprise Users: 300+



Tencent



Lenovo



Trip.com



Mobvista



PERFECT
WORLD

Industry-leading MPP Query Engine

01

Speed

- Industry leading query performance
- **Without de-norm'd tables**
- Supports aggregate or ad-hoc queries

03

Concurrency

- External analytics made possible
- **10,000+ QPS (Queries Pers Second)**
- Critical capabilities for modern apps

02

Real-time

- Events to insight in seconds
- Support high-speed ingestion
- **Updates and deletes in real time**

04

Operations

- Simple architecture low maint. effort
- Support for local disk storage or remote object storage
- **Unlimited horizontal auto scale**

Open Table
Formats (Read)



HIVE

ICEBERG




Apache

hudi




DELTA LAKE

Architecture




Different Data Architecture Models

Scatter/Gather




Single-Node Aggregation

Map Reduce




Barrier between stages
Staging data stored on disk

MPP




Pipelined Execution without barrier
Staging data at memory

Different Data Architecture Models




 SingleStore

 druid






 StarRocks


 trino

 ClickHouse


Key Features of StarRocks




Vectorized Query Engine



Multiple data storage
type support
(Local disk and Data
Lake)




Fully support
JOINS and
WIDE FLAT
TABLES




Real time
upsert and
merging from
batch and
stream

Performance!



StarRocks is **1.7X** faster than ClickHouse,
2.2X faster than Apache Druid.
SSB Benchmark Testing



StarRocks native is **14.6X** faster than Trino.
StarRocks Hive external table is **3.3X** faster
than Trino.

TPC-H Benchmark Testing

Summary: StarRocks makes Real-time Analytics Easy

Other Products

De-normalized table is a necessary evil

Struggle with updates and deletes

Low concurrency with only 10 – 100 users

Complex architecture, 3rd-party dependency, hard to maintain and scale



StarRocks

Superior query performance without de-normalization

Maintains performance while data is frequently updated/deleted

High concurrency with 10000s of users

Simplified architecture, easy to scale, and reduced TCO


Use Case: Tableau Dashboard at Airbnb

Background and Pain Points


The Airbnb Tableau Dashboard project is designed to serve both internal and external users by providing interactive dashboards. It requires a quick response to user queries. However, the query latency of previous solutions is over 10 mins, which is not acceptable. This project was just suspended until StarRocks is adopted.

StarRocks Solution

- StarRocks can directly connect and works very well with Tableau
- 3 tables (0.5B rows, 6B rows, 100M rows) + 4 joins + 3 distinct count + JSON functions and regex at same time, response time just 3.6s
- Reduce the query response time from mins level to sub-seconds level



Use Case: Minerva, the Metrics Store at Airbnb



Background and Pain Points

Minerva is Airbnb's internal unified metric platform

Vision: "Define metrics once, use them everywhere"

Over 30,000 metrics and 4,000 dimensions

Used for various data consumption scenarios, such as A/B testing, data exploration, and data analysis.

In Minerva v1, multiple flat tables to feed into Druid

Any change from source requires to refresh the wide tables, and it can take hours


Use Case: Trust Analytics at Airbnb

Background and Pain Points

To enhance security, Airbnb needs a real-time fraud detection system (Trust Analytics) to identify various attacks and take actions ASAP. This system must support Ad-Hoc query and real-time update.

StarRocks Solution

- StarRocks hosts real-time updated datasets via Primary Key.
- Dataset import from Kafka has a sub-minute delay.
- StarRocks provides second-level query latency for complex joins.
- Alerting can be achieved by just running a SQL query regularly.




User Case: Tencent IEG

Background and Pain Points

- . 400+ game data analysis and user behavior analysis
- . Operation reports need to be real-time.
- . Using ClickHouse for real-time analysis and Trino for Ad-hoc before, but they want to integrate them all.
- . Using Iceberg + COS store, need better performance.
- . Need elastic in ad-hoc query to deduce cost.

StarRocks Solution

- . Using StarRocks Primary key to solve update problem.
- . Using compute node on k8s to auto-scaling.
- . Get much more performance in ad-hoc query.



Use Cases: StarRocks Sweet Spot

OLAP multi-dimensional analytics

- User behavior analysis
- User profiling, label analysis, user tagging
- High-dimensional metrics report
- Self-service dashboard
- Service anomaly probing and analysis
- Cross-theme analysis
- Financial data analysis
- System monitoring analysis

Real-time analytics

- Online promotion analysis
- Logistics tracking and analysis
- Performance analysis and metrics computation for the financial industry
- Quality analysis for livestreaming
- Ad placement analysis
- Cockpit management
- Application Performance Management (APM)

High-concurrency analytics

- Advertiser report analysis
- Channel analysis for the retail industry
- User-facing analysis for SaaS
- Multi-tabbed dashboard analysis

Thanks for listening in on StarRocks

Get the Presentation and Links for more information on StarRocks

<https://bit.ly/starrocks-redhat>



Join our StarRocks Slack Community!

<https://bit.ly/starrocks-slack>



Red Hat
OpenShift



Try out our free tier on CelerData Cloud!

<http://cloud.celerdata.com>

Deploy our StarRocks Helm Chart on Red Hat OpenShift;
tutorial at

<https://github.com/StarRocks/starrocks/discussions/22767>

Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://www.linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://www.youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://www.facebook.com/redhatinc)



twitter.com/RedHat