



# Build User-Facing Analytics Application That Scales

Using StarRocks, a Linux Foundation Project

Albert Wong, [albert.wong@celerdata.com](mailto:albert.wong@celerdata.com)



StarRocks

# Agenda

- Trends and challenges in today's user-facing analytics
- How StarRocks solves the challenges of user-facing analytics
- Case studies of how StarRocks is being used for user-facing analytics



Profiting off of near real time knowledge - Battle of Waterloo, Nathan Rothschild and London Stock Exchange.

# Trends in user-facing analytics

User-facing analytics (UFA) is a rapidly growing field that is transforming the way businesses deliver insights to their users. UFA empowers users to explore and analyze data for themselves, without the need for technical expertise. This can lead to a number of benefits, such as:

1

Improved decision making

2

Increased user engagement


3

Reduced reliance on IT

Key trends in user-facing analytics:



Self-service Analytics



Embedded Analytics



Real-Time Analytics



Augmented Analytics




Conversational Analytics

# Trends in OLAP databases

Online analytical processing (OLAP) databases are evolving rapidly to meet the demands of modern data analytics. Here are some of the key trends in OLAP databases:

## 1 Cloud Native



## 2 A Sub-second vs. Second/Minute Query Response Time

## 2 B Streaming vs. Batch Data

## 2 C Mutable vs. Immutable Data

## 2 D Remote (Object) Storage vs. Local (SSD) Storage

## 2 E Open Table Format vs. Product Native Storage Format

## 3 Data Warehouse vs. Data Lake vs. Data Lakehouse



# Trends in OLAP databases

Open Lakehouse vs Proprietary / Hybrid Lakehouse

Proprietary / Hybrid

Open Storage

Open

Compute



Table Format



Storage Format





# Challenges in user-facing analytics

User-facing analytics (UFA) is a powerful tool that can help businesses deliver insights to their users, but it is not without its challenges. Here are some of the key challenges with UFA:

1

**Ingestion speed and data updates:**

User-facing analytics requires high-quality data that is available in real time. The logic and mathematics required to handle unbounded data sets is different than static.

2

**Query Response Time:** User-facing analytics needs to be able to deliver insights with sub-second response times versus seconds or minutes long query times.

3

**Scalability:** User-facing analytics solutions need to be scalable to handle large volumes of data and high concurrency. Performance and cost are always put to the test as a data set grows over time.

4

**Cost:** User-facing analytics solutions can be expensive to implement and maintain.




# How StarRocks solves the challenges of user-facing analytics





# StarRocks is an open-source query engine that delivers data warehouse performance on the data lake.


StarRocks enables:

- Directly query data on data lake and get data warehouse performance.
  - Sub-second joins and aggregations on billions of rows.
  - Serving hundreds of thousands of concurrent end-user requests.
  - No need for external data transformation tool for denormalization or pre-aggregation
  - Compatibility with standard sql protocol and Trino dialect.
- 
- Perform significantly more queries per second at lower latency on less hardware.
  - Increase the flexibility and capability of your data lake or analytics system.
  - Reduced cost and complexity by eliminating the need for costly data transformation pipelines.

# Low Latency Queries Solution


Sub-second query with thousands of QPS and beyond

## Scan



- Reducing IO
- SIMD filter

## Compression



- 4-8x compression rate
- Adaptive codec selection

## Indexes




~~min 100  
max 200~~



min 500  
max 800

- Min/max
- Ordinal index
- Secondary index

## Vectorized



- SIMD Instructions
- CPU prefetch and cache friendly
- Branch prediction friendly

**IO acceleration**

**CPU acceleration**



# High Concurrency

Sub-second query with thousands of QPS and beyond

## IO Bound Workload

- Less complex queries such as point queries.
- Very High QPS.
- Typically requires Latency in the low 10s of ms on massive amount of data (billions of rows).

## CPU Bound Workload

- Complex OLAP-style queries with JOINS and AGGs.
- Interactive analytics.
- Reports, dashboards, etc.



# High Concurrency Solution

## IO Bound Workload

Scan the least amount of data possible

1. Bucketing: Fine-grained control over how the data is distributed in your cluster.
  - Avoid data skew.
  - Choose a high cardinality column that often appears in your `WHERE` clauses.
2. Indexing: Use prefix index (order key), Bitmap, Bloom Filter index to minimize disk access.
3. IO-optimized Hardware:
  - Configure (multiple) high-performance SSD drives.
  - Good network in terms of latency and throughput.

## CPU Bound Workload

Pre-computation or reusing previous results to take load off your CPUs

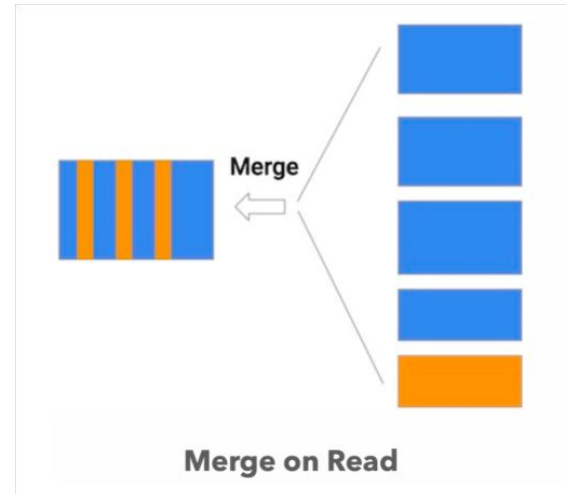
1. Intelligent Caching:
  - Final Result Cache: For the folks who keep hitting the same button just for fun.
  - Query Cache (Intermediate Result Cache): Reuse partial results even if the queries are not identical.
1. Pre-computation:
  - Denormalization through Column Partial Update.
  - Pre-aggregation through MVs and Aggregate Table.
  - Generated Column (3.1): Materialize a column.



# Fresh Mutable Data

Mutable data is important!!

Existing solutions relies on merge on read – Huge compromise on query performance.

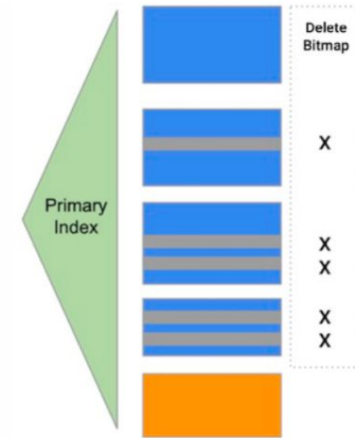


# Fresh Mutable Data Solution

Delete and insert mechanism with Primary Key index.

Support Real-Time mutable data with no compromise on query performance.

Approach	Pros	Cons
Delete and insert	Simple, efficient for read-heavy workloads	Can be inefficient for write-heavy workloads
Merge on read	Can be efficient for write-heavy workloads	More complex than delete and insert



**Delete and Insert**



# Resource Isolation

Why Resource Isolation is Needed:

- Maintaining Business-Critical Operations: Uninterrupted and undisturbed performance for key tasks.
- Quality of Service Assurance: Resource assurance for varying user groups in a multi-tenant environment.

Current Approach: Physical Isolation:

- Pros:
  - Effective isolation of resources.
- Cons:
  - Low Hardware Utilization: Results in underutilized resources due to lack of sharing or overprovisioning.
  - High Costs: Need to meet peak demand of each user group leads to idle resources and escalated costs, especially in large-scale operations.
  - Gets worse when you grow to tens of thousands of tenants.



# Resource Isolation Solution

Resource Group: Support over-provisioning CPU resources.

Three types of workloads defined:

- Short query: Time sensitive queries with dedicated resource (no over-provisioning).
- Query queue: Queries that are not time sensitive but are critical (cannot be killed).
- The rest: supports customized rules to kill big queries.




**StarRocks is an open-source query engine that delivers data warehouse performance on the data lake.**



**Reduce the time and cost of developing data analytics projects.**

- No ingestion and data copying.
- Stop paying to denormalize data that may never be queried.
- Keep the tools you've been using.




**Sub-second query while serve millions of users.**

- Multi-dimensional interactive analytics through on-the-fly computations.
- Single source of truth on open data lake analytics, with no external system for caching or pre-computation.



**Make real-time decisions in all business scenarios, especially when updated data, like in logistics, is needed.**


- Real-time update without sacrificing query performance.
- Simpler real-time data pipeline
- Ditch stateful stream processing jobs (denormalization & preaggregation) with efficient on-the-fly computation.



**Data warehouse query performance on the data lakehouse with no data copying:** Natively integrates with open data lake including Hive, Hudi, Iceberg, and Delta Lake.

**Ditch Denormalization:** Perform joins on multiple tables with millions of rows in seconds.

**No need for external data transformation tool:** Perform on-demand pre-computation (like denormalization) within StarRocks, eliminating another processing tool in your data pipeline.



#### Intelligent Query Planning

- Cost-based optimizer generates optimized query plan.
- Global runtime filter.

#### Efficient query execution

- In-memory data shuffling enables fast and scalable JOIN operation.
- C++ SIMD-optimized columnar storage and vectorized query executions deliver the industry's fastest query performance.

#### High concurrency

- Built-in Materialized View.
- An Intelligent caching system.
- Secondary Indexes.



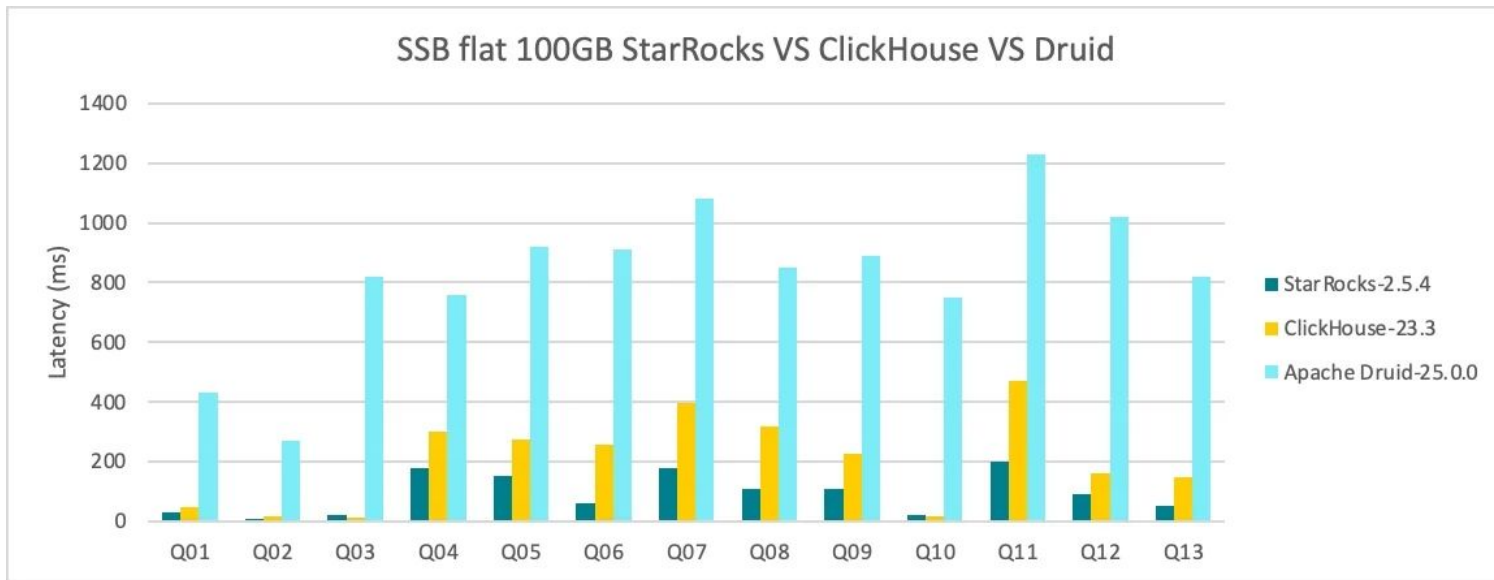
**Ditch stream processing tools for denormalization:** StarRocks' multi-table performance allows you to ditch the rigid and expensive stateful stream processing jobs for denormalization and pre-aggregation.

**Real-time updates:** Through primary key table, support real-time data updates while having no impact on query performance.

**Real-time Query Performance at scale:** The synchronized materialized view (MV) further accelerates aggregated queries at scale for real-time analytics.

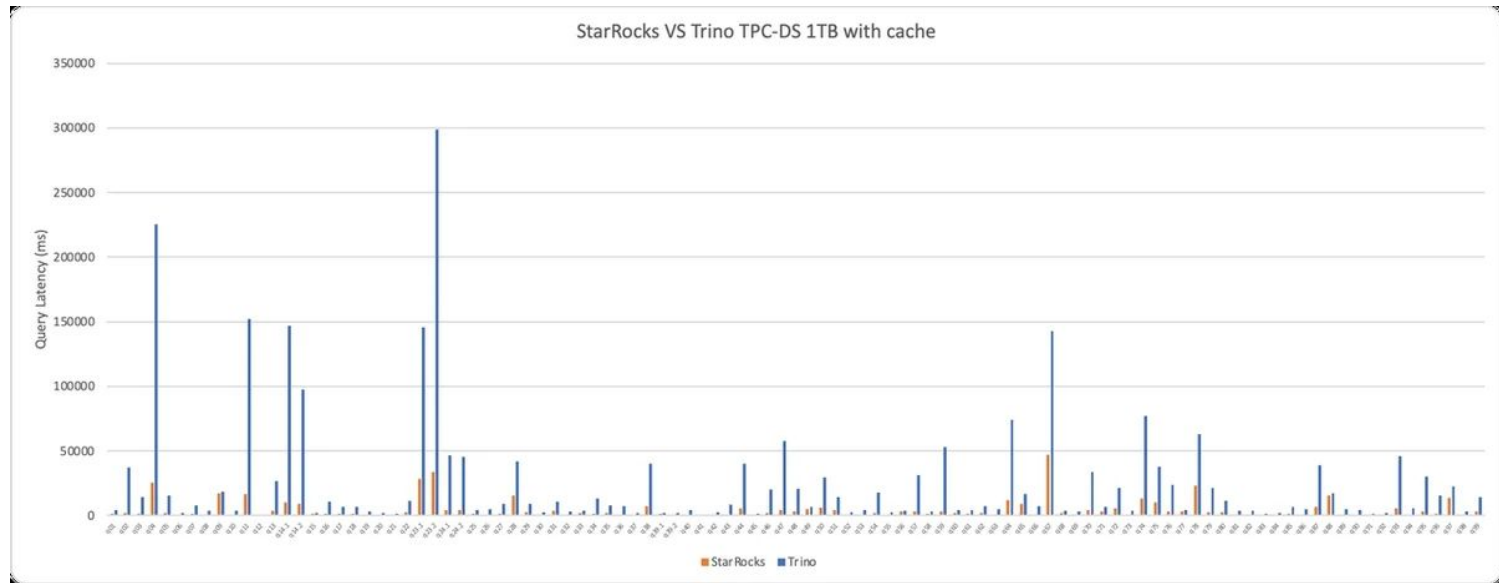
# Benchmark

StarRocks Offers 2.2x Performance over ClickHouse and 8.9x Performance over Apache Druid® in Wide-table Scenarios Out of the Box using product native table format.



# Benchmark

StarRocks Delivers 5.54x Query Performance over Trino in Multi-table Scenarios using Apache Iceberg table format with Parquet files.

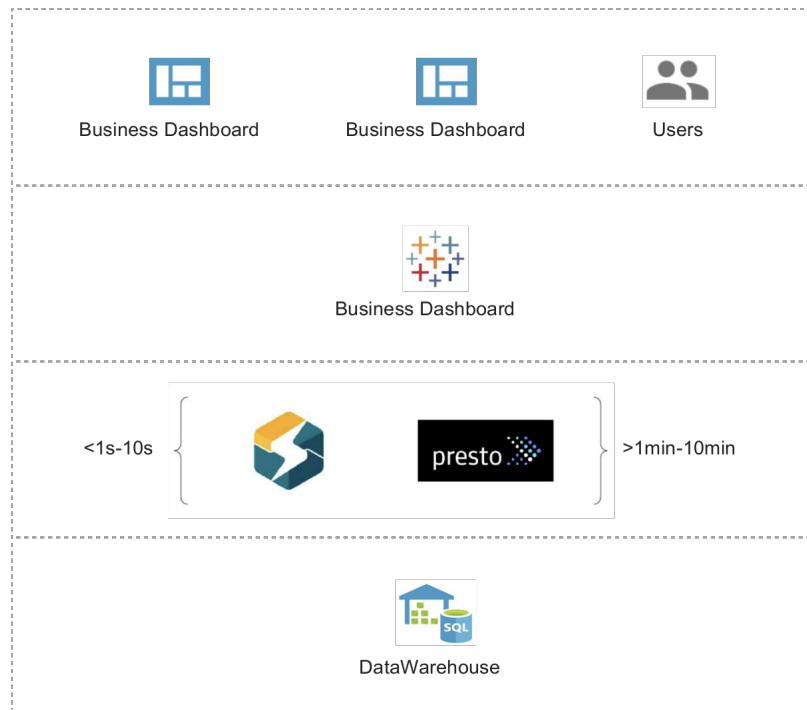


# Use Case: Tableau Dashboard at Airbnb

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.

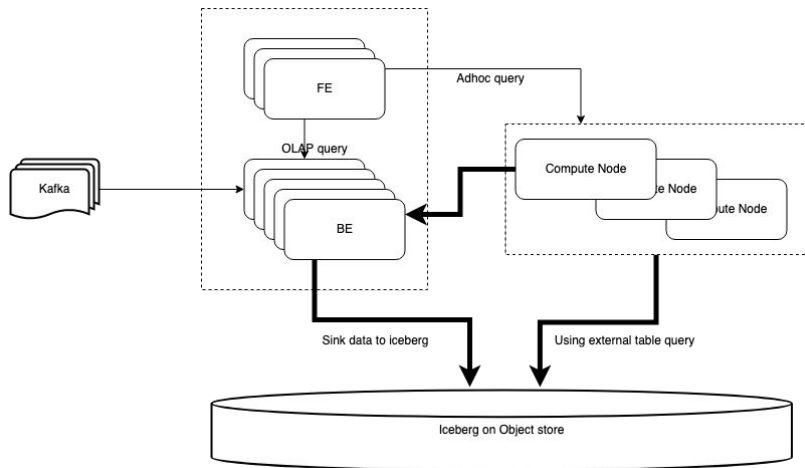



# Use Case: Game and User Behavior Analytics at Tencent IEG

- 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.

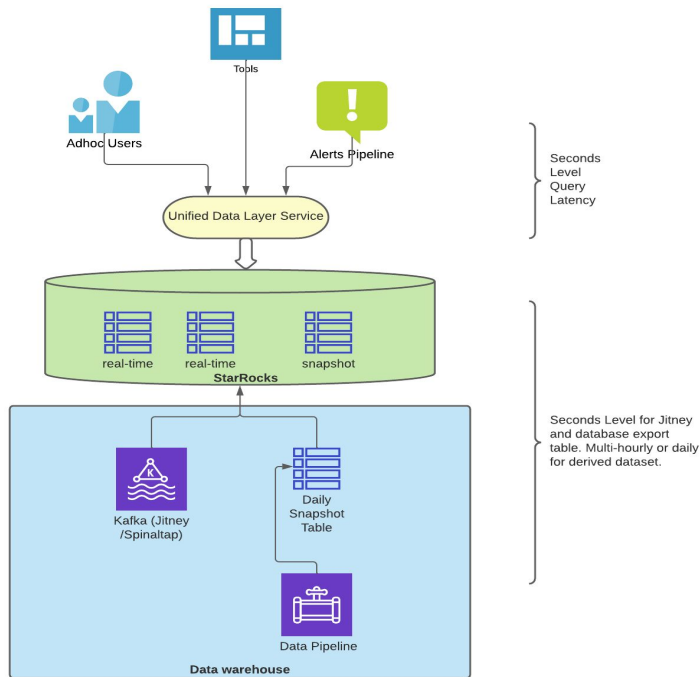


# Use Case: Trust Analytics at Airbnb

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.

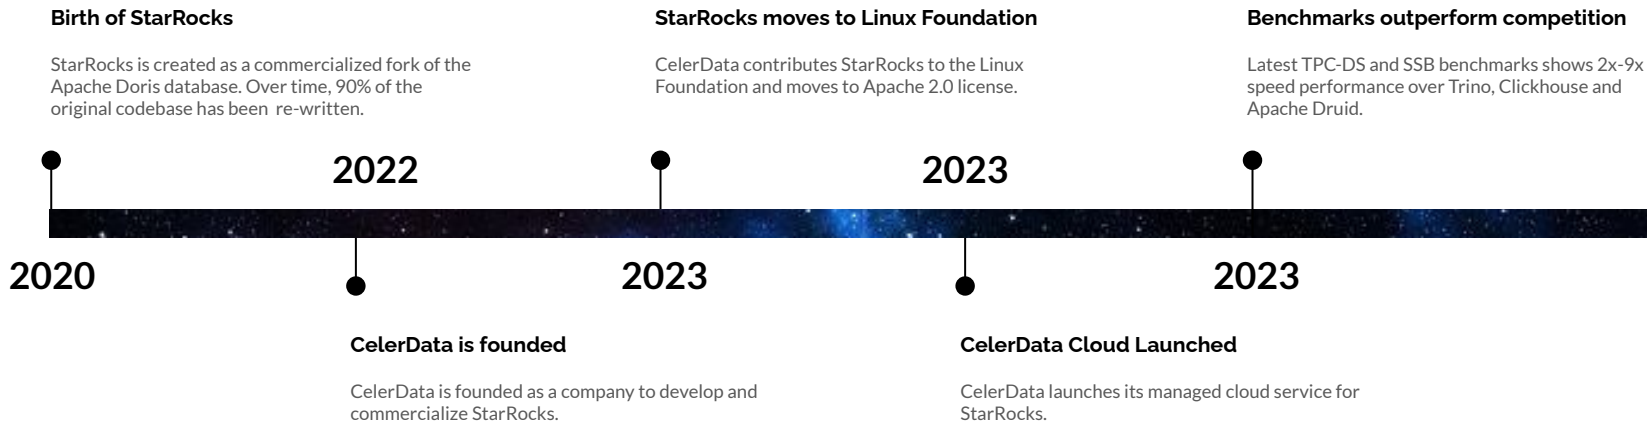






# History of StarRocks and CelerData

StarRocks was designed to address the challenges of real-time analytics, including the need to support high concurrency, low latency, and a wide range of analytical workloads. StarRocks also offers a number of features that are not available in other real-time analytics databases, such as the ability to query data directly from data lakes.



# Thank you.



• Website	<a href="http://starrocks.io">starrocks.io</a>
• Managed Service	<a href="http://cloud.celerdata.com">cloud.celerdata.com</a>