

Making HTAP a reality

With Postgres (+ lakehouse)

Zhou Sun
Mooncake Labs

HTAP: Still the Dream, a Decade Later



Dani Palma

Follow

10 min read · Jun 18, 2025

▲ HTAP is Dead (mooncake.dev)

159 points by moonikakiss 3 months ago | hide | past | favorite | 90 comments

Musings

HTAP is Dead

👤 Zhou Sun 📅 May 4, 2025 ⌚ 5 min read

rip htap.



Why Hybrid Transactional-Analytical Processing Remains an Organizational Challenge



Convex

Subscribe

convex

database

ANDREW PAVLO

Associate Professor of Computer Science, Carnegie Mellon University


Full screen (f)

Shorts




mooncake

Why HTAP?




Rowstore



`SELECT * FROM trades WHERE id=123;`

Fast for OLTP, Slow on OLAP

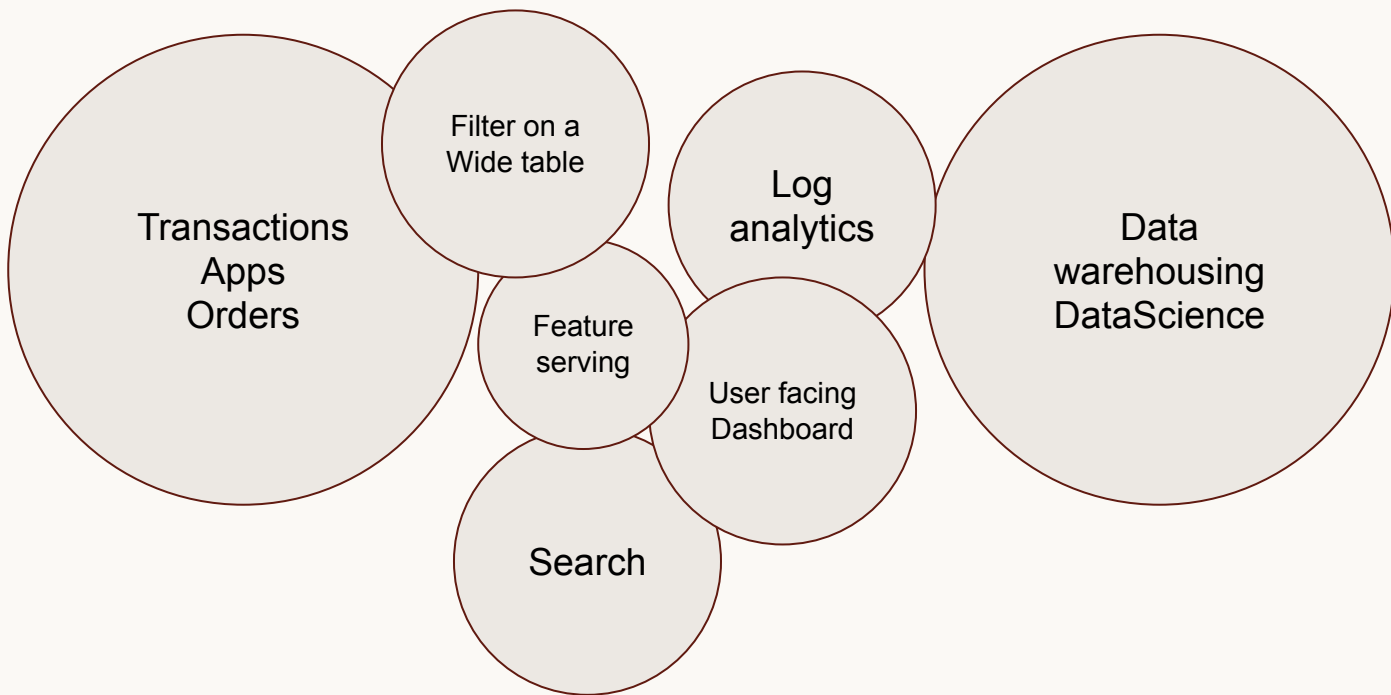
Columnstore



`SELECT symbol, count(*) FROM trades
GROUP BY 1 ORDER BY 2 DESC;`

Fast for OLAP, Slow on OLTP

The Reality?



What went wrong?

OLTP ++

Oracle/ sql-server

Great, Reliable

And \$\$\$\$

HTAP feature:

- Cluster columnstore key
- In memory columnstore
- HeatWave

No real analytics

- Single box
- Vendor Lock

Hook two proprietary systems

Oracle/ sql-server

A new OLTP engine

- **You may not trust as much**

TiKV

Modified Mysql

TiDB / ByteHTAP

A new OLAP engine

- **Won't evolve as fast**

Diverged Clickhouse 2018

Flink

HTAP native engine

Oracle/ sql-server

Cool, Performant

But need sacrifice one or both

TiDB / ByteHTAP

SingleStore / HANA/ CedarDB

I want the best

OLTP

DB2



MySQL



NoSQL (MongoDB)



NewSQL (Spanner, Cockroach, Vitess)



Postgres 

OLAP

DB2



Vertica



Hadoop




cloud data warehouse



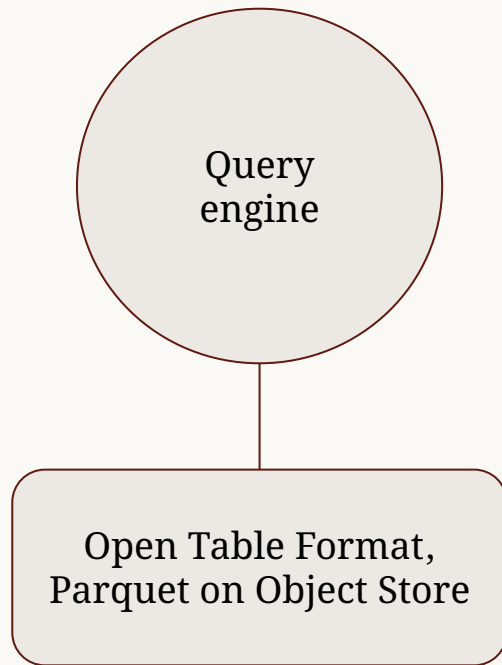

LakeHouse 


1980s


2020s




What is a Lakehouse









Mooncake is a composable framework


1. Replicate from Postgres To Lakehouse

2. Fast analytics queries against Lakehouse

Logical replication just works...

- Streaming Xact
- 2PC
- Parallel apply

And for Analytics



What if


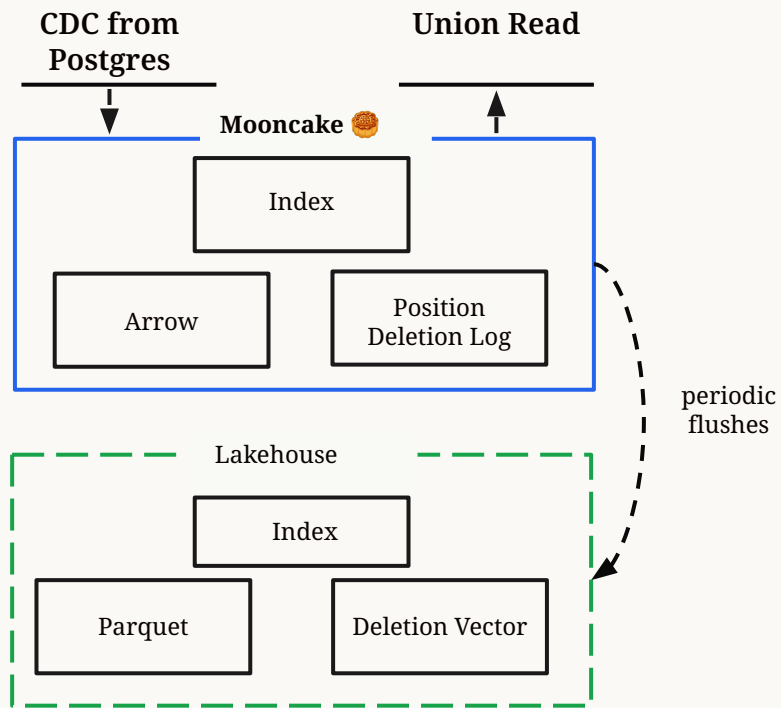
Build an engine from scratch

To handle writes from PG?

Zero-ETL

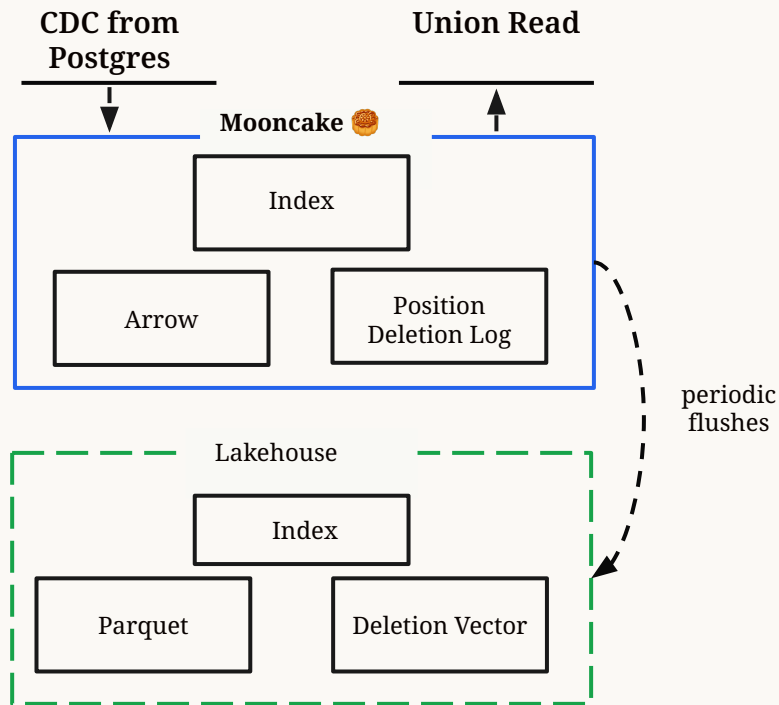
!= Write Some Logs to Object-store

- Columnstore
- Up-to-Date
- Queryable
- Optimized



Real-time Columnstore

- Buffer
 - No small parquet files
- Index
 - No delay for update/deletes
- Write-aware table optimization
 - No conflicts




Postgres-native

- Streaming Transaction
 - Apply transaction immediately & Commit or Rollback
 - **Pick any existing CDC tool for any system**
 - Try delete a 100M row table from postgres
 - See the delay
- KeepAlive LSN & Exactly once
- Initial Loading

Fast Query

Mooncake/Lake from Postgres



Vectorized Execution on Columnstore

Processes data in small batches rather than row-by-row

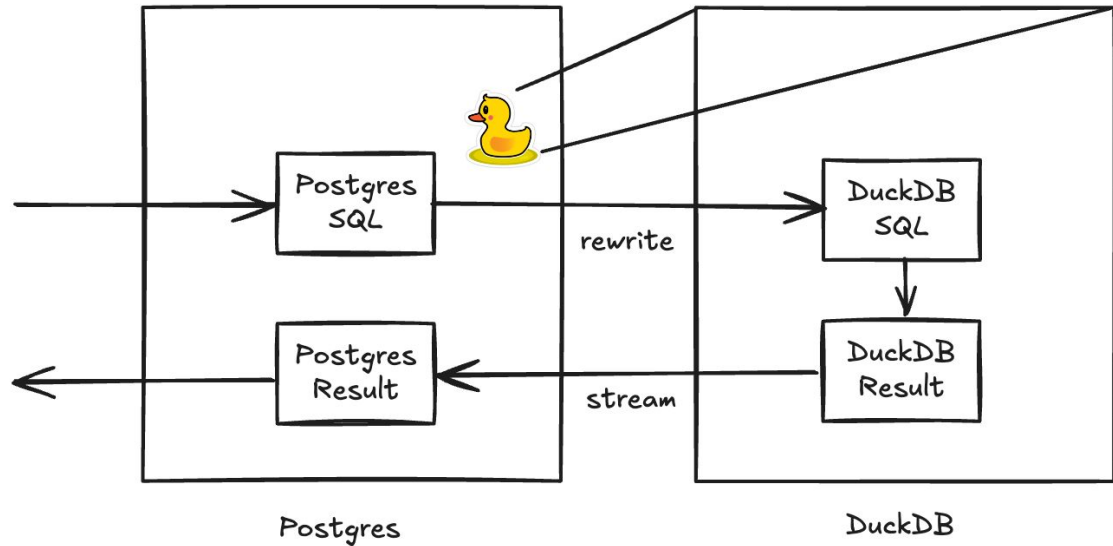
- SIMD acceleration
- Encoded data
- Minimize function calling overhead

dictionary	
AAPL	0
AMD	1
AMZN	2

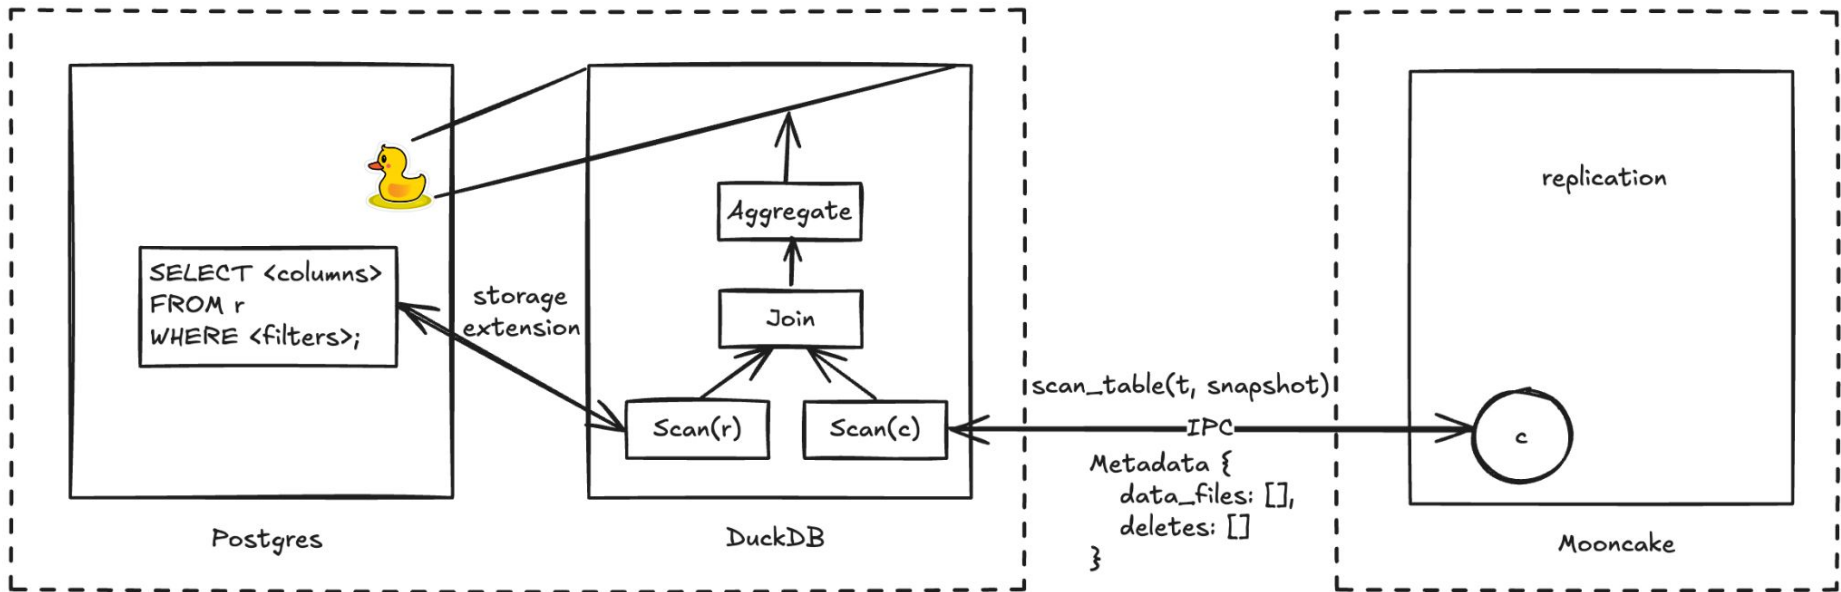
	symbol
AMD	1
AAPL	0
AAPL	0
AMZN	2
AMD	1
AMZN	2
AAPL	0

```
SELECT symbol, count(*)  
FROM trades GROUP BY 1;
```

pg_duckdb



pg_mooncake

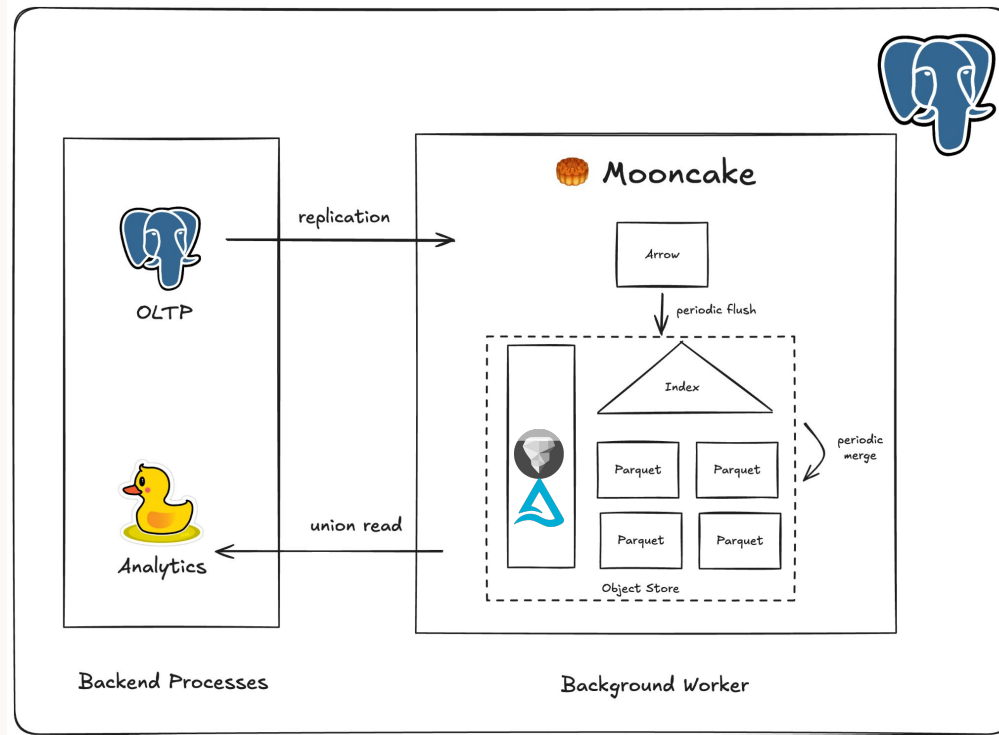


Read optimizations

- Consistent Read
 - Read with LSN, Mooncake wait until that LSN
- Union Read
 - LakeHouse Files (Parquet + DeletionVector)
 - Local Files (Arrow + New Deletion Log)

Start from

A Postgres Extension




And Scale to Infinite




OLTP

Read replica



replication



union read

Backend Processes

 Mooncake

Arrow

periodic flush

Index

Parquet

Parquet

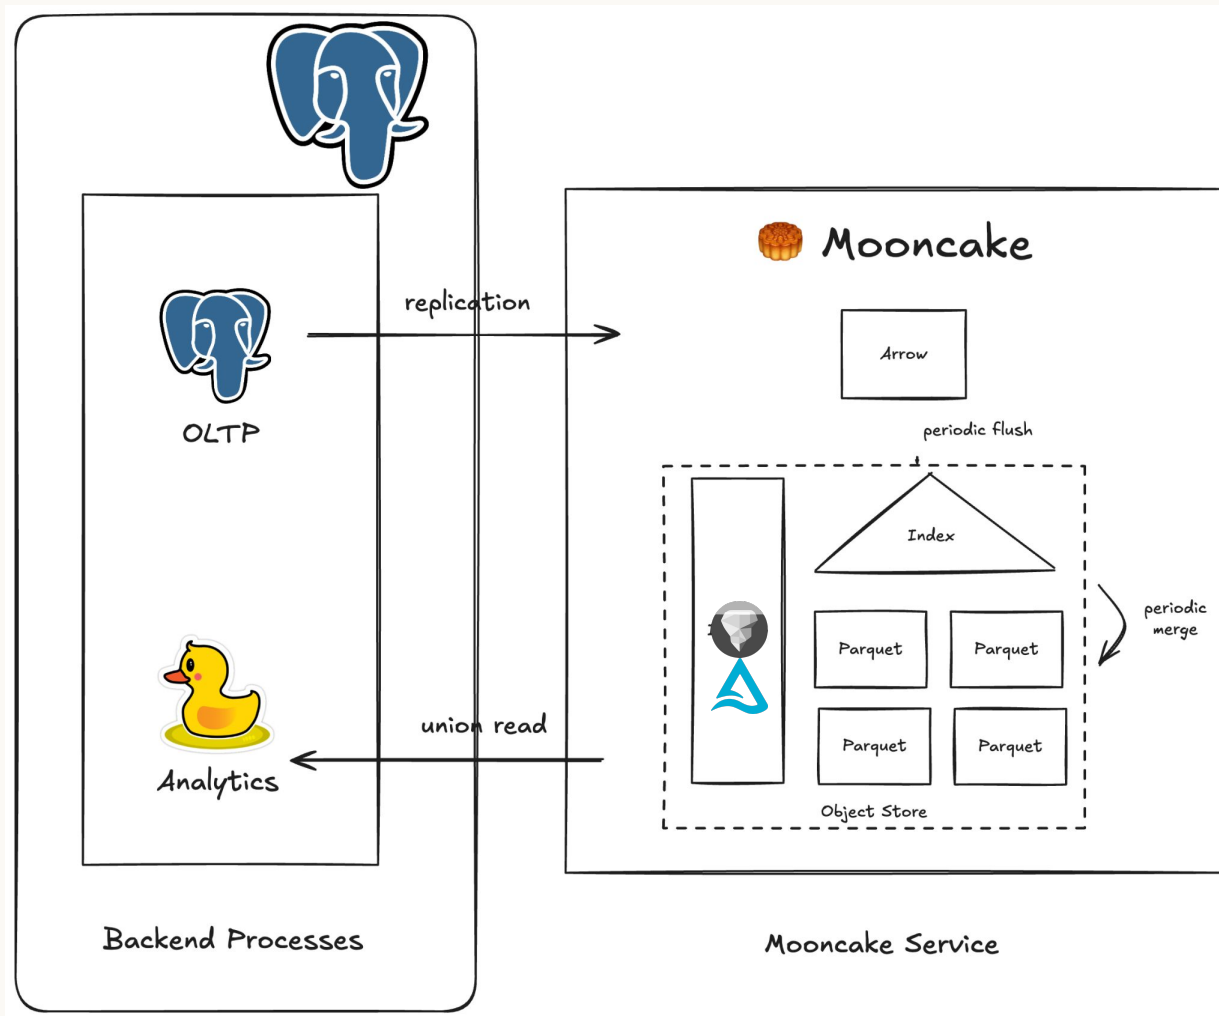
Parquet

Parquet

periodic merge

Object Store

Background Worker





OLTP

replication



Analytics

union read



Mooncake

Arrow

periodic flush

Index

Parquet

Parquet

Parquet

Parquet

Object Store

periodic merge

Mooncake Service

Spark



trino



Some numbers

Postgres Just Cracked the Top Fastest Databases for Analytics (mooncake.dev)

378 points by moonikakiss 6 months ago | [hide](#) | [past](#) | [favorite](#) | [120 comments](#)

Ch-Bench (TPC-C + analytics)

Collocated pg_mooncake, On I4I.4X , 1000 Warehouses, 300MB/s WAL

Replication Delay < 1s , and analytics freshness < 1.5s

Analytics queries: Avg < 10s

Official benchmark releasing soon

Musings

HTAP is Dead Simple

👤 Zhou Sun 📅 May 4, 2025 ⌚ 5 min read

rip htap.



Thank you!

- Website: <https://mooncake.dev/>
- Email: founders@mooncakelabs.com
- pg_mooncake: https://github.com/Mooncake-Labs/pg_mooncake (MIT License)
- moonlink: <https://github.com/Mooncake-Labs/moonlink>