

OSA CON 24



# UNIFIED DATA MANAGEMENT WITH CLICKHOUSE<sup>®</sup> AND POSTGRES

---



**Shivji Kumar Jha**  
Staff Engineer at Nutanix



**Sachidananda Maharana**  
MTS IV at Nutanix

November 19-21, 2024 | Free, online.

# About The Speakers



Shivji Kumar Jha  
Data Platform Engineer

- Software Engineer & Regular Speaker / Meetups
- Excited about:
  - Distributed Databases & Streaming
  - Open-Source Software & Communities
  - MySQL/Postgres, Pulsar/NATS, Druid/Clickhouse



Sachidananda Maharana  
Software Engineer  
OLAP Ninja

- Regular Platform Engineer
- Excited about:
  - Distributed OLAP Databases
  - Open-Source Enthusiast



# Legal Disclaimer

## Forward Looking Statements

This presentation and the accompanying oral commentary may contain express and implied forward-looking statements, including, but not limited to, statements relating to: our business plans, initiatives and objectives; ability to execute such plans, initiatives and objectives in a timely manner, and the benefits and impact of such plans, initiatives and objectives, including our ability to manage our expenses in future periods; our financial targets and our plans to achieve those targets; the benefits and capabilities of our platform, products, services and technology; our plans and expectations regarding new products, services, product features and technology including those that are still under development or in process; the timing of any product releases or upgrades or announcements; anticipated trends, growth rates and challenges in our business and in the markets in which we operate; our ability to develop new solutions, product features and technology and bring them to market in a timely manner, as well as the impact and/or benefits of including additional solutions or features in our product portfolio; market acceptance of new technology and recently introduced solutions; the interoperability and availability of our solutions with and on third-party platforms, including public cloud platforms; our ability to maintain and strengthen our relationships with our channel partners, OEMs and other third parties, and the impact of any changes to such relationships on our business, operations and financial results; the competitive market, including our competitive position and ability to compete effectively, our projections about our market share in future periods, the competitiveness of our future cost structure with those of other companies, and the competitive advantages of our products; our plans and timing for, and the success and impact of, our transition to a subscription-based business model; and macroeconomic trends and geopolitical environment, including the on-going global supply chain disruptions.

These forward-looking statements are not historical facts and instead are based on our current expectations, estimates, opinions, and beliefs. Forward-looking statements should not be considered as guarantees or predictions of future events. Consequently, you should not rely on these forward-looking statements. The accuracy of these forward-looking statements depends upon future events and involves risks, uncertainties, and other factors, including factors that may be beyond our control, that may cause these statements to be inaccurate and cause our actual results, performance or achievements to differ materially and adversely from those anticipated or implied by such statements, including, among others, the risks detailed in our most recent Annual Report on Form 10-K and Quarterly Reports on Form 10-Q, each as filed with the U.S. Securities and Exchange Commission, the SEC, which should be read in conjunction with the information in this presentation and accompanying oral commentary. Our SEC filings are available on the Investor Relations section of our website at [ir.nutanix.com](http://ir.nutanix.com); the SEC's website at [www.sec.gov](http://www.sec.gov). These forward-looking statements speak only as of the date of this presentation and accompanying oral commentary and, except as required by law, we assume no obligation, and expressly disclaim any obligation, to update, alter or otherwise revise any of these forward-looking statements to reflect actual results or subsequent events or circumstances.

## Product or Roadmap Information

Any future product or roadmap information included in this presentation and the accompanying oral commentary is (i) intended to outline general product directions, (ii) not a commitment, promise or legal obligation for Nutanix to deliver any material, code, or functionality, and (iii) not intended to be, and shall not be deemed to be, incorporated into any contract. This information should not be used when making a purchasing decision. We note that Nutanix has made no determination as to whether separate fees will be charged for any future products, product enhancements and/or functionality which may ultimately be made available. Nutanix may, in its discretion, choose to charge separate fees for the delivery of any future products, product enhancements and/or functionality which are ultimately made available.

## Third Party Reports and Publications

Certain information contained in the presentation and accompanying oral commentary made available as part of this digital event may relate to or be based on reports, studies, publications, surveys and other data obtained from third-party sources and our own internal estimates and research. While we believe these third-party reports, studies, publications, surveys and other data are reliable as of the date of the applicable presentation, they have not been independently verified, and we make no representation as to the adequacy, fairness, accuracy, or completeness of any information obtained from third-party sources.

## Trademark Disclaimer

© 2022 Nutanix, Inc. All rights reserved. Nutanix, the Nutanix logo, and all Nutanix product, feature, and service names mentioned herein are registered trademarks or trademarks of Nutanix, Inc. in the United States and other countries. Other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Nutanix may not be associated with, or be sponsored or endorsed by, any such holder(s).

# CONTENTS

- Background
  - The explosion of data
  - Variety of data & query patterns
- Running an open-source data stack YOURSELF
  - Operational challenges and cost
  - Distributed system is hard in the real world!
- More pieces in data stack => more complexity
- Thought Experiment: (Postgres, Clickhouse)
  - Is that enough?
  - Does it work well together?





**So much data & So many databases. How do I map?**



# Volume, Variety, and Use Cases

- Complexity from data diversity
  - Volume: Data size grows exponentially
    - Compared to 2010 we are generating 76x data per year in 2024 . (details in upcoming slides)
  - Variety: Different types of data (structured, semi-structured, unstructured)
    - Row oriented
    - Column oriented
    - Documents(JSON, Protobuf) – text, audio, video
    - Graph data
    - Vector Embeddings?
  - Use-Cases: Expanding range of query patterns and applications across business functions
    - Transactions
    - Join Queries
    - Filters and Window Functions
    - Realtime or Batch or both (Lambda, Kappa?)
    - Slice and dice queries: group by, Sorting and Limiting, Aggregations etc
    - Lossy, Lossless, performance, dollars, availability etc etc





Other top of mind concerns...

ACTIVITY

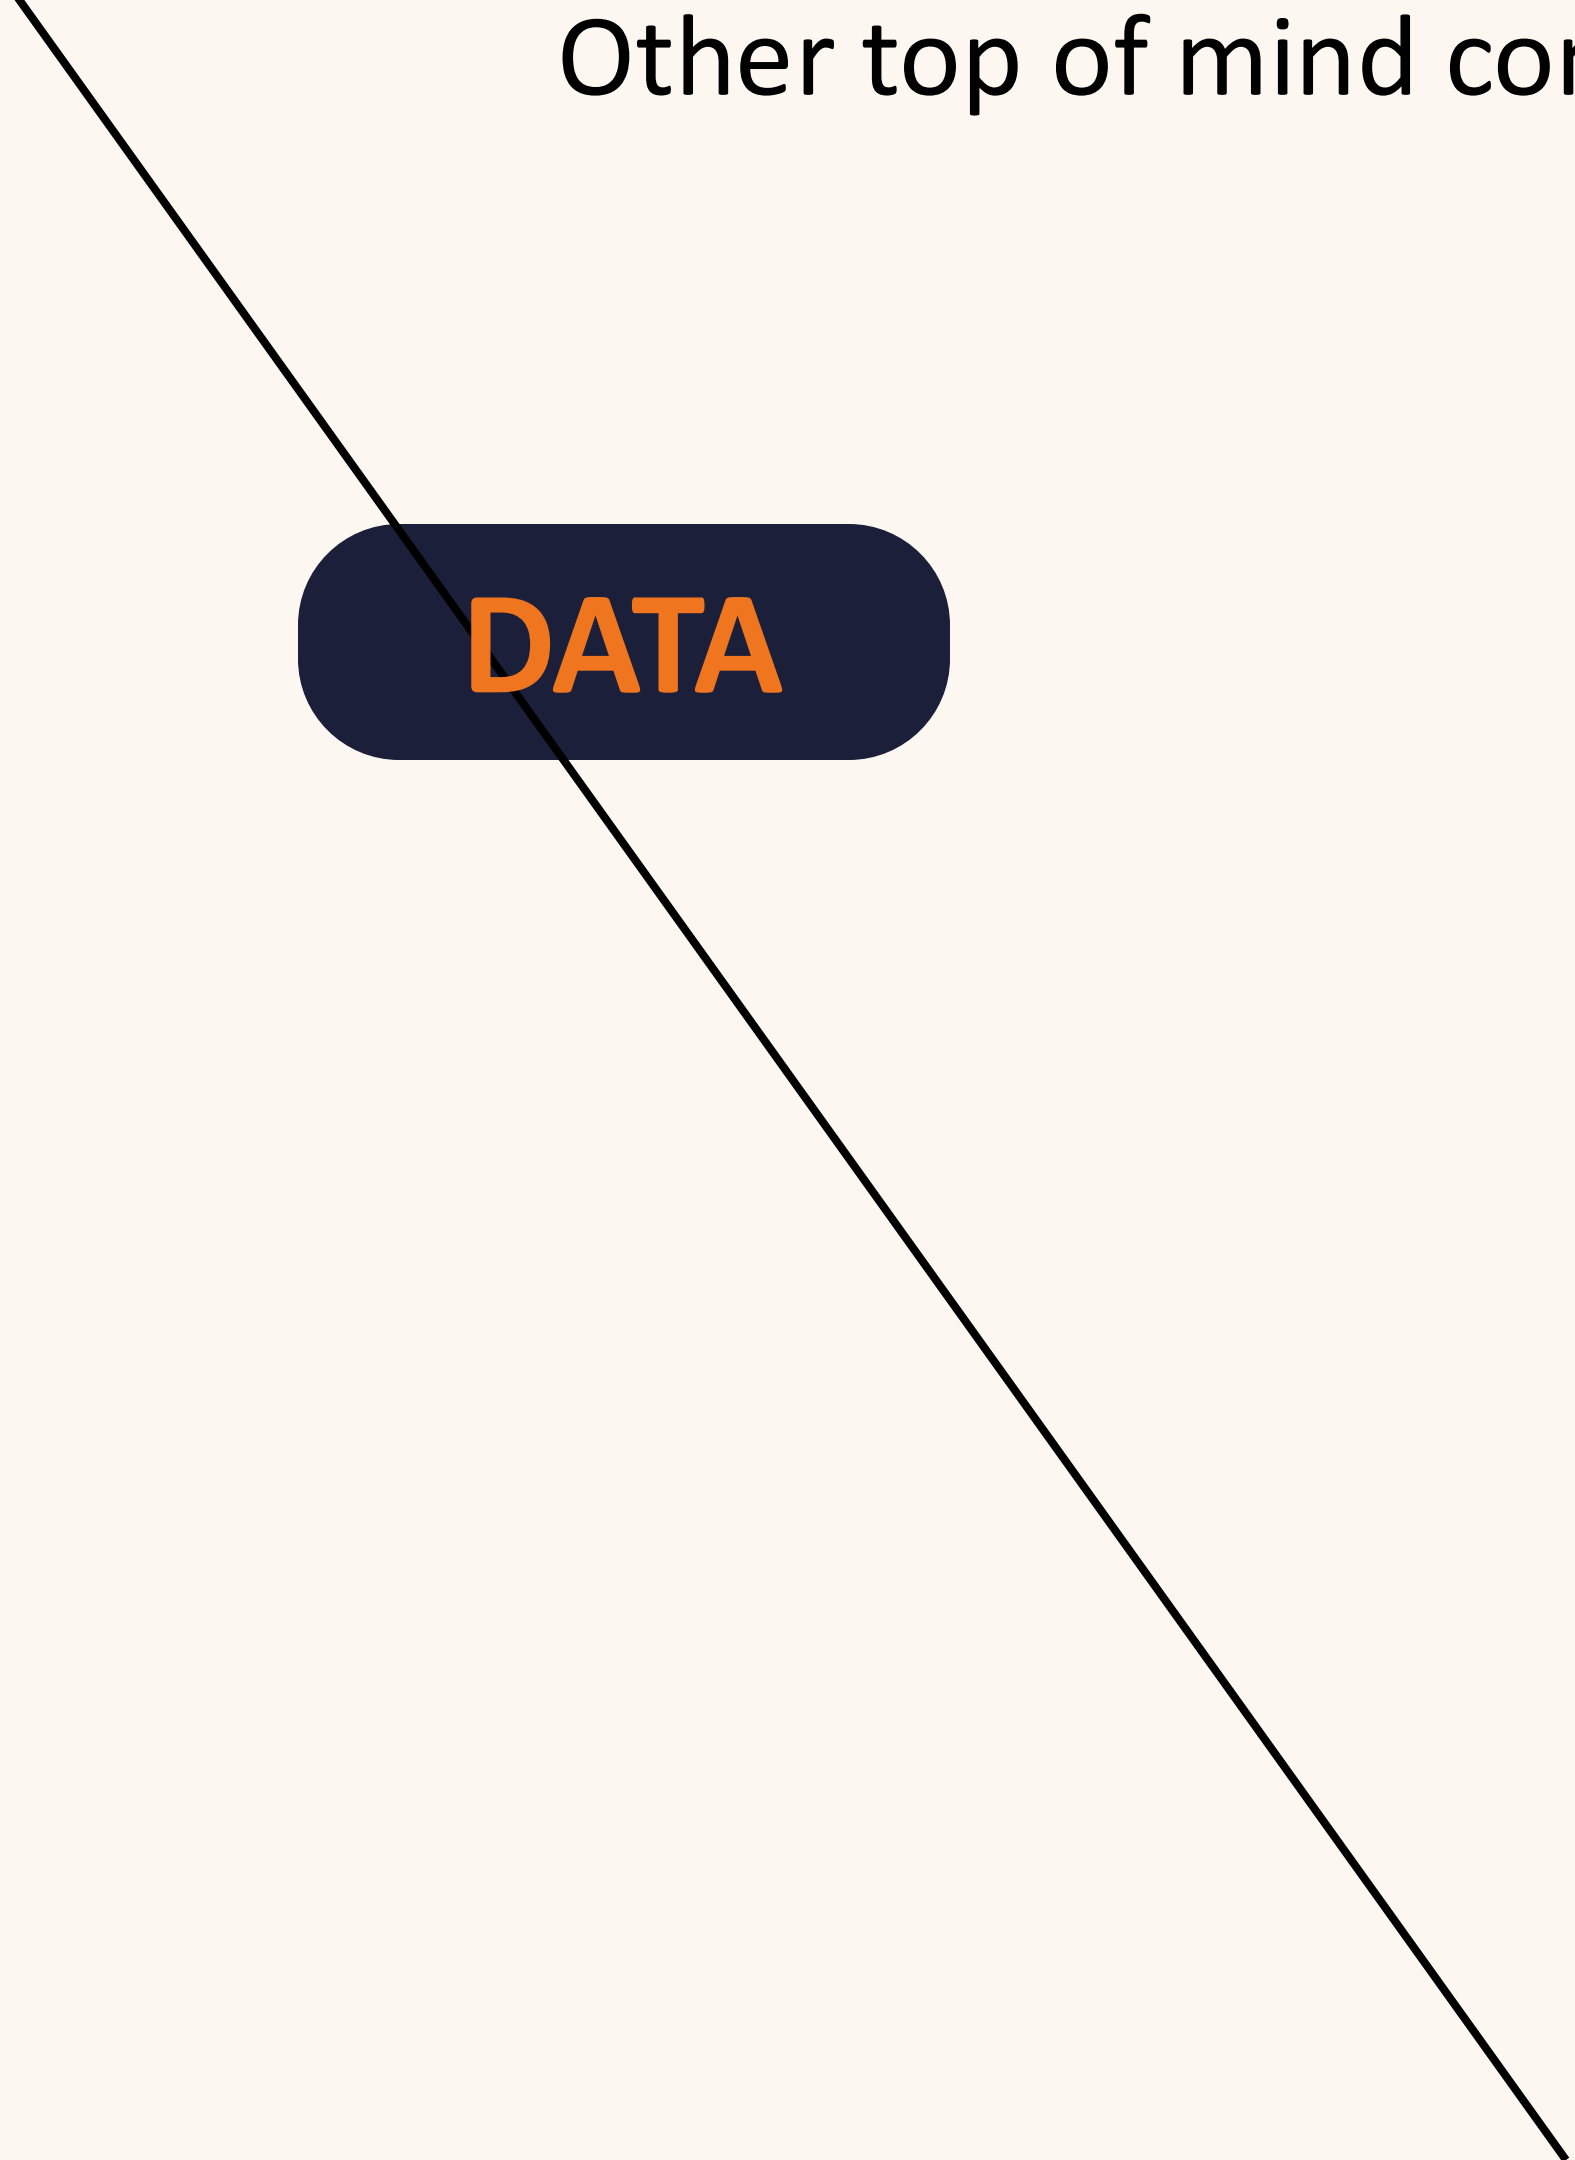
TRANSACTIONS

ANALYTICS

METRICS

LOGGING

**DATA**



Other top of mind concerns...

ACTIVITY

TRANSACTIONS

ANALYTICS

METRICS

LOGGING

**DATA**

AI HYPE

AM I LOSING ANY DATA?

HOW WILL I STORE IT?

HOW WILL I SAVE COST?

HOW WILL I SECURE IT?



# The Data Explosion

TRANSACTIONS

DATA


VECTOR DB

ANALYTICS

METRICS

LOGGING

Year	Data generated
2010	2 zb (x)
2015	15.5 zb (8x)
2020	64 zb (32x)
2024*	147 zb (74x)





**Financial**

**TRANSACTIONS**

**ANALYTICS**

**User interaction**

**User Activity**

**DATA**

**METRICS**

**Time Series**

**AI ML**

**VECTOR DB**

**LOGGING**

**Historical logs**



Oracle  
SQL Server  
MySQL

Financial

TRANSACTIONS

ANALYTICS

User interaction  
User Activity

Snowflake  
Druid / Pinot  
DuckDB

DATA

METRICS

Time Series

Prometheus  
TimeScaleDB

Pinecone  
Milvus/QuestDB

AI ML

VECTOR DB

LOGGING

Historical logs

ElasticSearch  
Splunk



<https://dbdb.io>

Discover and learn about 1,014 database management systems

Oracle  
SQL Server  
MySQL

Pinecone  
Milvus/QuestDB

Financial

AI ML

TRANSACTIONS

VECTOR DB

ANALYTICS

DATA

METRICS

LOGGING

User interaction

Time Series

Historical logs

User Activity

Snowflake

Prometheus

ElasticSearch

Druid / Pinot

TimeScaleDB

Splunk

DuckDB

# Challenges of Running A Database in Enterprise Environments

- If the DBaaS mode works for you, that's a fine option
  - Cost affordable at scale?
  - Vendor lock-in?
  - Flexibility?
- Running yourself has complex challenges that go beyond just adding resources. Managing these systems requires
  - specialized expertise
  - often necessitating a dedicated platform team for configurations,
  - resilience testing, and legal compliance approvals.
  - For example, large enterprises must consider factors like system performance, high availability, and load balancing across nodes to ensure uptime and meet demands, especially in mission-critical applications.





# Deeper Technical Challenges

---

**1. Compliance and Security:** GDPR, HIPAA, region-specific storage, encryption etc

**2. Reliability and Uptime:**

1. High availability with replication, backup, and disaster recovery

**3. Performance Optimization:**

1. Data and access patterns evolve
2. Load balancing, caching, and efficient query handling requires ongoing monitoring and configuration updates


**4. Resource and Cost Management:**

1. Scaling resources dynamically in cloud environments costly.
2. Running on-prem opens up need for more specialization.

This operational complexity demonstrates the importance of a skilled platform team and highlights why many enterprises consider managed database services or carefully select database technologies with robust enterprise support.

# Challenges in Distributed Systems (Auto-scale)


- Operating distributed systems is complex
  - Coordination, consistency, and fault tolerance are difficult to achieve
    - Leader election
    - Ensure linearizability
  - Complexities in scaling maintaining these systems
    - Patch update
    - Zero downtime upgrade
    - Seamless migration
- Impact on performance and reliability when not handled well
  - Prod down time = loss of business and trust
  - Cost
  - Data loss






# Many Distributed Databases = Exponentially More Complexity

- More databases, more problems
  - Managing multiple databases in a distributed environment increases the complexity exponentially
    - Cost
    - Duplicate data
    - More Experts required
    - Cloud and onPrem
  - Challenges with data consistency, replication, and cross-database querying
    - Data spread across database
    - Different queries
  - Increased operational overhead for managing and monitoring
    - Patch
    - Upgrade
      - Infra
      - new release



# Simplifying Complexity - Reducing Moving Parts


- Consolidate where possible
  - Reduce number of systems to reduce complexity
    - Save cost, disk, expertise, support
    - Less duplicate data
  - Aim for solutions that balance flexibility with simplicity
    - Monitor at one place
    - Less Experts needed
  - Leverage multipurpose databases like Postgres and Clickhouse for versatility
  - Keep options open
    - A good architecture is one that commits least (OSS)







We Love  
PostgreSQL



# Why Postgres?

- Advantages of using PostgreSQL
  - As open as you get them!
  - Mature, robust, and highly extensible
  - Supports a wide variety of data types and use cases
  - Strong community and ecosystem support
  - Proven reliability for both transactional and analytical workloads



"[PostgreSQL](https://db-engines.com/en/blog_post/106) is our DBMS of the Year for the fourth time, after winning already three titles between 2017 and 2020."

[https://db-engines.com/en/blog\\_post/106](https://db-engines.com/en/blog_post/106)





# Postgres supercharged with Extensions


- Postgres is the most popular database
- **And there is so much innovation happening with its extensibility**
  - **Hydra, pg\_duckdb, citus** etc for analytics
    - **Parade\_db** : Postgres for Search and Analytics
    - **Postgres for logs**: parseable
    - **Postgres for timeseries**: timescale
    - **Postgres on S3** : NeonDB
    - Postgres for Vector embeddings : pg\_vector

## **A piece of opinion..**

- If I start a company today and have little data, this works like a charm!!!
- But these extensions are not mature enough for enterprise use-cases YET!
- Neither are they mature enough to handle huge scales.



We Love  
Clickhouse





# Clickhouse: A Rising Star in Analytics

- High-performance analytics with Clickhouse
  - Columnar storage optimized for analytical queries
  - Exceptional performance for large-scale data analysis
  - Increasingly popular for real-time analytics and time-series data
  - Complements PostgreSQL for high-speed data querying
- But not for “transactions”!
- Hates updates/deletes, well (almost)!

# Lots of buzz around clickhouse lately



The Cloudflare Blog

Subscribe

Email Address

- Product News
- Speed & Reliability
- Security
- Zero Trust
- Developers
- AI
- Policy
- Partners

## Log analytics using ClickHouse

2022-09-02

### ClickHouse is in the house

Insights gained and lessons learned from our long video analytics migration journey.



zeev · [Follow](#)

Published in Vimeo Engineering Blog · 12 min read · Dec 12, 2023

229 5



Uber Blog

[Sign up](#)

Engineering

## Fast and Reliable Schema-Agnostic Log Analytics Platform

19 February 2021 / Global

### Druid Deprecation and ClickHouse Adoption at Lyft



Ritesh Varyani · [Follow](#)

Published in Lyft Engineering · 12 min read · Nov 29, 2023

549 6





# Postgres and Clickhouse

## “together”

---





# Postgres & Clickhouse “together” !

- Powerful combination for diverse data needs
  - Use Postgres for complex transactional data and flexible schema support
  - Use Clickhouse for high-speed analytics and heavy query loads
  - Together, they offer a powerful blend of reliability and speed for modern data requirements





# Data Movement Between PostgreSQL and Clickhouse

- **Key Point:** Seamless **bi-directional data movement** between PostgreSQL and Clickhouse enables versatile use cases, allowing teams to combine strengths of both databases.
- **Tools**
  - [PostgreSQL Table Engine](#)
    - Query postgres data directly from clickhouse
  - [Altinity clickhouse connector](#)
    - Moves data in real time from transactional database tables in MySQL and PostgreSQL to ClickHouse for analysis.
  - [Clickhouse fdw](#)
    - Use the clickhouse\_fdw extension provided by Supabase to start querying data in ClickHouse from within Postgres. Variety use cases
      - Consolidating to a single query interface to support both transactional (OLTP) and analytical (OLAP) applications
      - Providing a simple way to join transactional data stored in Postgres with analytical data such as logs, telemetry, events, and more.



# PostgreSQL and Clickhouse: Use cases & Success Stories!

## Use Cases:

- **Hybrid Analytical Workloads:** Run real-time analytics in Clickhouse, then sync processed results back to PostgreSQL for integration into transactional applications.
- **High-Performance Reporting:** Transfer large datasets to Clickhouse for intensive analytical queries while PostgreSQL handles transactional processing.
- **Data Archiving and Query Offloading:** Archive older data from PostgreSQL to Clickhouse to reduce storage and query load while retaining fast access to historical data.

## Success Stories:

- **eCommerce Platforms:** Companies use Clickhouse for detailed customer behavior analysis and then sync relevant results to PostgreSQL for integration with customer management systems.
- **Financial Services:** Banks utilize this setup for real-time trading analytics in Clickhouse, syncing essential results to PostgreSQL for transactional integrity.
- **IoT Data Processing:** IoT providers analyze large volumes of sensor data in Clickhouse and aggregate results in PostgreSQL for reporting and device management.

This integration gives companies the flexibility to build robust and scalable architectures using PostgreSQL and Clickhouse, supporting diverse, data-intensive applications.

# PostgreSQL Table Engine

- In postgres
  - `CREATE TABLE leads (id INTEGER PRIMARY KEY, name VARCHAR);`
- In Clickhouse
  - `CREATE TABLE psql_table`  
(  
`id` Int32,  
`name` String,  
)  
`ENGINE = PostgreSQL('localhost:5432','test','leads', 'postgres','pwd');`

- The PostgreSQL engine allows SELECT and INSERT queries on data stored on a remote PostgreSQL server.
- Contd.



# PostgreSQL Table Engine Demo

```
clickhouse
2024.11.16 17:08:21.234539 [ 116353 ] {} <Debug> DNSResolver: Updating DNS cache
2024.11.16 17:08:21.234612 [ 116353 ] {} <Debug> DNSResolver: Updated DNS cache

psql
test=# select * from leads; Corresponding PostgreSQL Table 1
 id | name
----+-----
  1 | x
  2 | y
  5 | z
  6 | w
  7 | p
(5 rows)

test=# exit
sachidananda.m@VVHXTGQJKT:~$ psql -U postgres
Password for user postgres:
psql (14.13 (Homebrew))
Type "help" for help.

postgres=# \c test
You are now connected to database "test" as user "postgres".
test=# \dt
      List of relations
 Schema | Name  | Type  | Owner
-----+-----+-----+-----
 public | leads | table | postgres
(1 row)

test=# select * from leads; Data Inserted into Postgres Table 4
 id | name
----+-----
  1 | x
  2 | y
  5 | z
  6 | w
  7 | p
  8 | shiv
(6 rows)

test=# |

clickhouse
of simultaneous queries.

VVHXTGQJKT :) show tables

SHOW TABLES
Query id: d4f5eacb-baea-4bc9-bc92-1582e057fa1c
1. psql_table PostgreSQL Table Engine 2
2. tp

2 rows in set. Elapsed: 0.001 sec.

VVHXTGQJKT :) select * from psql_table;

SELECT *
FROM psql_table
Query id: 1faeb7d3-c7a0-4b92-82cc-ab27a631ce59
 id | name
-- | ---
1.  1 | x
2.  2 | y
3.  5 | z
4.  6 | w
5.  7 | p

5 rows in set. Elapsed: 0.019 sec.

VVHXTGQJKT :) insert into psql_table values(8,'shiv') Insert into Clickhouse Table 3

INSERT INTO psql_table FORMAT Values
Query id: 14761237-00ed-43e7-bdbe-7a06fc1fcfae
Ok.

1 row in set. Elapsed: 0.004 sec.

VVHXTGQJKT :) |
```

---

# The Amazing Duo (Postgres + Clickhouse)

← Back

[Blog / Company and culture](#)

## ClickHouse acquires PeerDB to boost real-time analytics with Postgres CDC integration



ClickHouse Team

Jul 30, 2024 - 4 minutes read

“PeerDB’s technology significantly reduces the time it takes to connect data in transactional databases like Postgres with analytical databases like ClickHouse, boosting developer productivity and time-to-value when building real-time data driven applications. With PeerDB CDC technology integrated into ClickHouse Cloud, users will benefit from seamless, real-time data replication from Postgres databases to ClickHouse.”

# Efficient Data Transfer with PeerDB between PostgreSQL and ClickHouse

## Overview of PeerDB:

- Real-time replication from PostgreSQL to ClickHouse,
- Change Data Capture (CDC) to track data updates.
- Low-latency sync without impacting source DB performance.

## How It Works:

- Uses PostgreSQL Logical Replication to capture changes and transfer them to ClickHouse with minimal lag.
- Data is batched and sent in real time, reducing transfer times (e.g., from hours to minutes for initial loads).
- Supports schema evolution, row-level transformations (e.g., masking PII), and robust data consistency checks, ensuring ClickHouse holds accurate and up-to-date data.



Here is the Picture we  
wanted to  
**PAINT!**

---

# The Amazing Duo (Postgres + Clickhouse)

- Reduce number of databases
- Build a Platform team with concentrated expertise
- Deep in-house experience in leveraging these technologies
  - Skilled in implementing, tuning, and optimizing Postgres and Clickhouse
  - Deep understanding of distributed data challenges and solutions
    - Mitigate crash scenarios
  - Our small and focussed team ensures seamless integrations and performance optimization

**Postgres**

**Financial**

**TRANSACTIONS**

**ANALYTICS**

**User interaction**

**User Activity**

**Postgres  
Extensions**

**Clickhouse**

**DATA**

**METRICS**

**Time Series**

**Postgres  
Extensions**

**Clickhouse**

**pg\_vector**

**AI ML**

**VECTOR DB**

**LOGGING**

**Historical logs**

**Clickhouse**



Open for

Q n A

Thank You

Contact us @

Shivji Kumar Jha - <https://www.linkedin.com/in/shivjiha/>

Sachidananda Maharana - <https://www.linkedin.com/in/sachidanandamaharana/>

