

# TimescaleDB 2.0

A decorative graphic consisting of a network of nodes and connections. The nodes are represented by three stacked circles, and the connections are represented by dotted lines. The network is spread across the top half of the slide, with nodes positioned around the text 'TimescaleDB 2.0'.

**Time-series данные  
в распределенном кластере TimescaleDB  
поверх ОПСУБД PostgreSQL**

# Спикер



Иван Муратов  
CTO @ WalIoT  
<https://github.com/binakot>

KRASNODAR  
DEV DAYS

KRASNODAR  
BACKEND

Первая  
Мониторинговая  
Компания


WALIOT



PostGIS



TimescaleDB

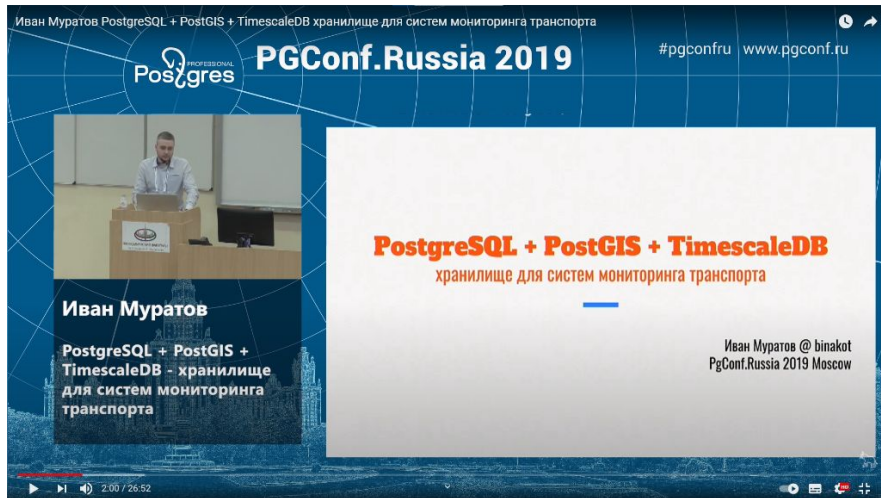


PostgreSQL

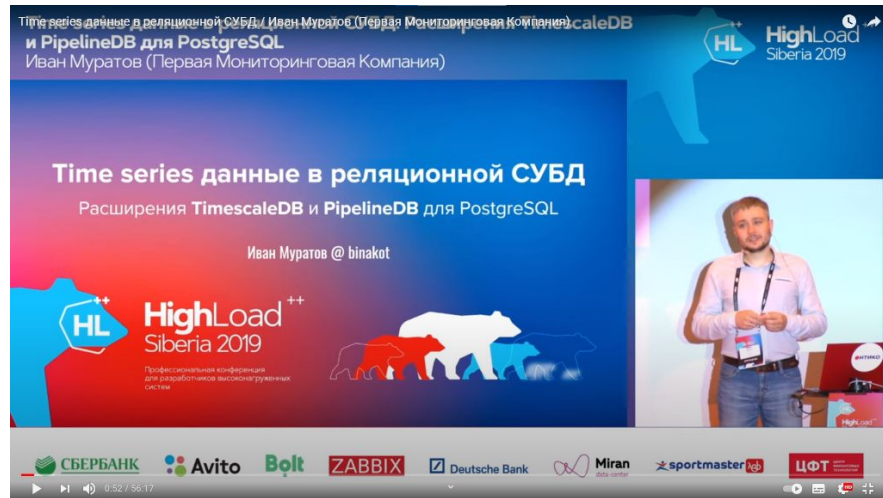


WALIOT 

# В предыдущих сериях...




<https://youtu.be/SEncf-h4Npw>




<https://youtu.be/3WkNp7mllv0>

## Trend of the last 24 months



# DB-Engines Ranking of Time Series DBMS



# Time series

**Временные ряды** — собранные в разные моменты времени данные о характеристиках исследуемого процесса.

Временные ряды состоят из двух элементов:

1. временная метка;
2. характеристики, называемые уровнями ряда.

Особенности time-series данных:

- Time-centric
- Append-only
- Recent

# HET!



**postgresmen**  
@postgresmen

Following

## One SQL to Rule Them All: Management of Streams and Tables


[arxiv.org/abs/1905.12133](https://arxiv.org/abs/1905.12133)

▲ truth\_seeker 10 hours ago [-]

Great write-up. SQL is the most underrated language.  
I am doing most of the stuff mentioned in that paper right now in my project with PipelineDB extension.  
PostgreSQL 11 +  
CitusDB extension (horizontal scaling and sharding) +  
cstore\_fdw extension (columnar data storage) +  
file\_fdw (to treat file as table) +  
PipelineDB extension (Streaming computations and MV) +  
TimescaleDB extension (timeseries data storage)  
Also, PG 12 release will have many optimizations for table partitioning, query planning and its parallel execution

<https://news.ycombinator.com/item?id=20059006>






Software  
Engineers



**TIMESCALE**

# TimescaleDB



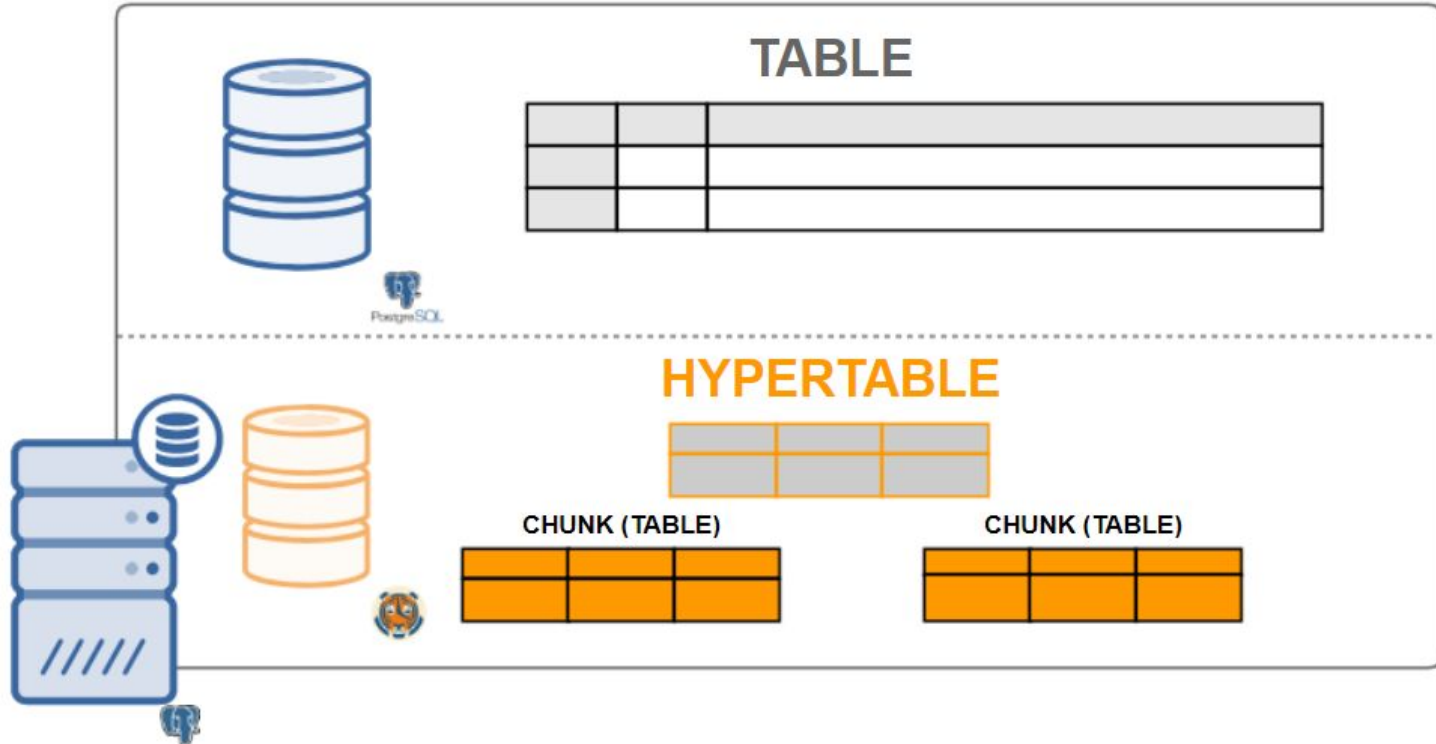
- Open source (Apache 2.0, Timescale License)
- Self-hosted / Cloud (AWS, Azure, GCP) / Fully-managed (SaaS)
- PostgreSQL ecosystem + TimescaleDB tools (Promscale, Tune, etc)
- Hypertables and distributed hypertables
- Full SQL + TimescaleDB hyperfunctions
- Real-time continuous aggregates
- Data retention
- Downsampling
- User-defined actions
- Native compression
- Massive scaling
- Function pipelines\*






# Single-Node Instance

# Single-Node Instance



# Hypertable of chunks

- Автоматическое секционирование/партиционирование по времени
- Уменьшение размеров секций/партиций (чанков) и индексов к ним
- Увеличение скорости вставки
- Упрощение обслуживания и счастье для DBA 

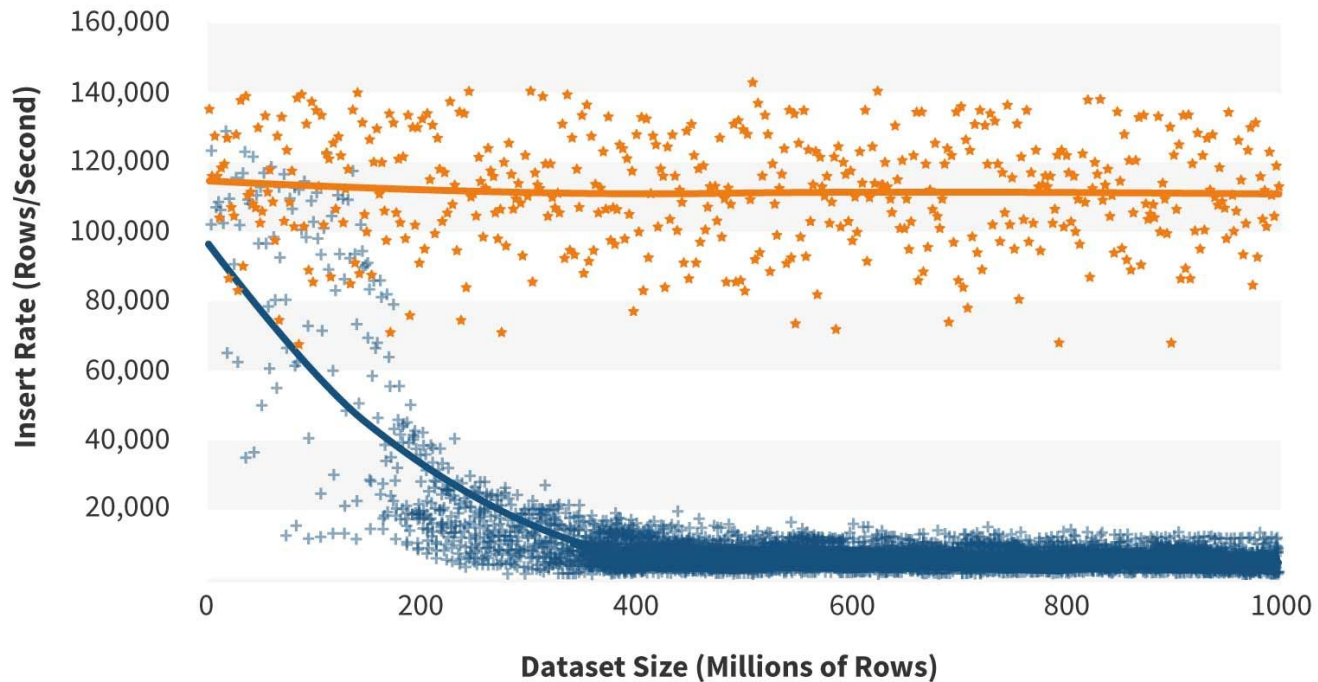
```
CREATE TABLE telemetries (  
    imei          TEXT          NOT NULL,  
    time         TIMESTAMPTZ   NOT NULL,  
    latitude     DOUBLE PRECISION NOT NULL,  
    longitude    DOUBLE PRECISION NOT NULL,  
    speed        SMALLINT      NOT NULL,  
    course       SMALLINT      NOT NULL,  
    CONSTRAINT telemetries_pkey PRIMARY KEY (imei, time)  
);
```

```
SELECT * FROM create_hypertable (  
    'telemetries', 'time',  
    chunk_time_interval => INTERVAL '7 days'  
);
```

```
INSERT INTO telemetries VALUES ...;
```

```
SELECT  
    time_bucket('30 days', time) AS bucket,  
    imei,  
    avg(speed) AS avg,  
    max(speed) AS max  
FROM telemetries  
WHERE speed > 0  
GROUP BY imei, bucket  
ORDER BY imei, bucket;
```

# Ingest Rate: PostgreSQL vs. Timescale



+ — PostgreSQL

Insert batch size: 10,000 Rows

Final avg. throughput: 5k (PG) vs. 111k (TS)

★ — TimescaleDB

Cache: 16 GB Memory

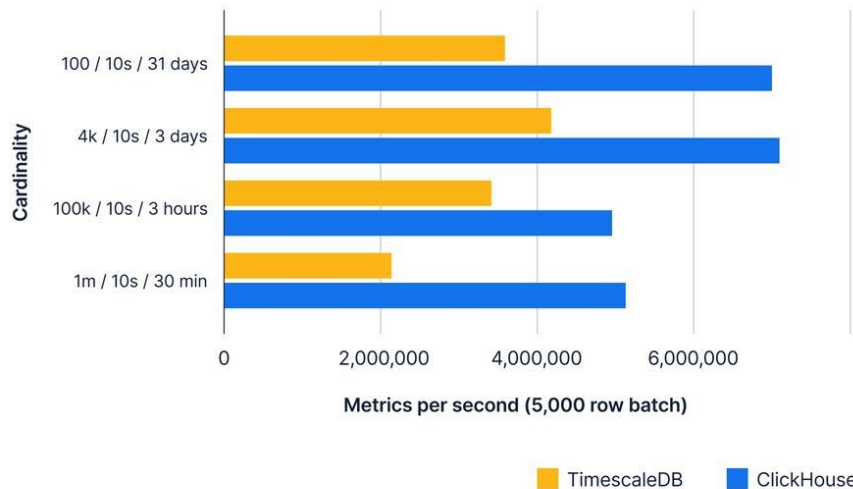
Time to load 1B rows: 37.9h (PG) vs. 2.6h (TS)



# Benchmarking, not “Benchmarketing”

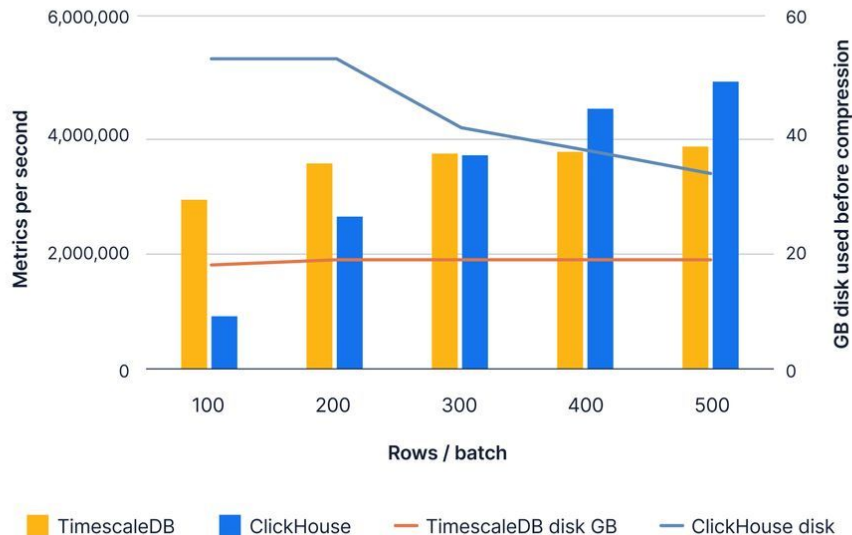
## TimescaleDB vs ClickHouse

Insert rate vs. cardinality



## TimescaleDB vs ClickHouse

Insert rate vs. Batch size





# Multi-Node Cluster

A top-down view of a desk with a light-colored wooden surface. In the center is the book 'Designing Distributed Systems' by Brendan Burns. To the top left is a small green plant in a black tray. To the right is a white coffee cup filled with black coffee. To the left is a rotary telephone with a tan handset and a silver base. The book cover features a blue header with the O'Reilly logo, a central illustration of two birds, a large blue title box, and a white footer with the author's name.

O'REILLY

A detailed black and white illustration of two birds, possibly sparrows, facing each other. The bird on the left is smaller and has a white patch on its throat. The bird on the right is larger and has a white patch on its cheek. They are perched on a thin branch.

# Designing Distributed Systems

PATTERNS AND PARADIGMS FOR SCALABLE, RELIABLE SERVICES

Brendan Burns



Microsoft Azure

# Clustering system for horizontal scaling

MySQL -> Vitess (Cloud Native Computing Foundation graduated)

PostgreSQL -> ?



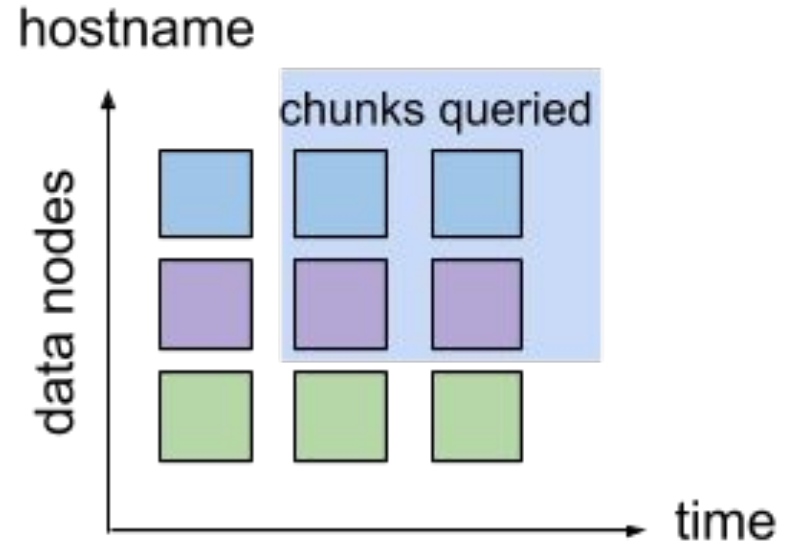
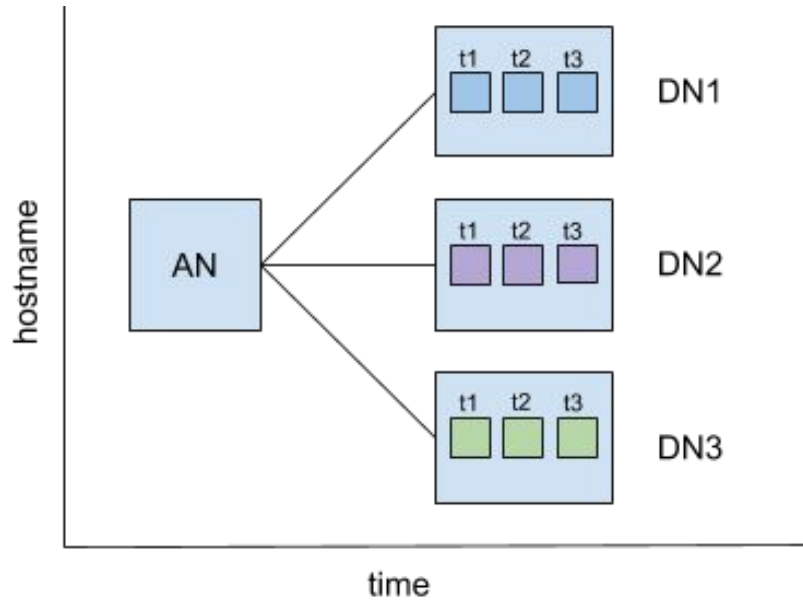
# Distributed hypertables

- Time-and-space шардирование без необходимости ребалансировки
- Узел доступа (Access Node) и множество узлов данных (Data Nodes)
- Распараллеливание запросов к DN и агрегация результатов на AN
- Нативная репликация (replication\_factor)

# Multi-Node Cluster

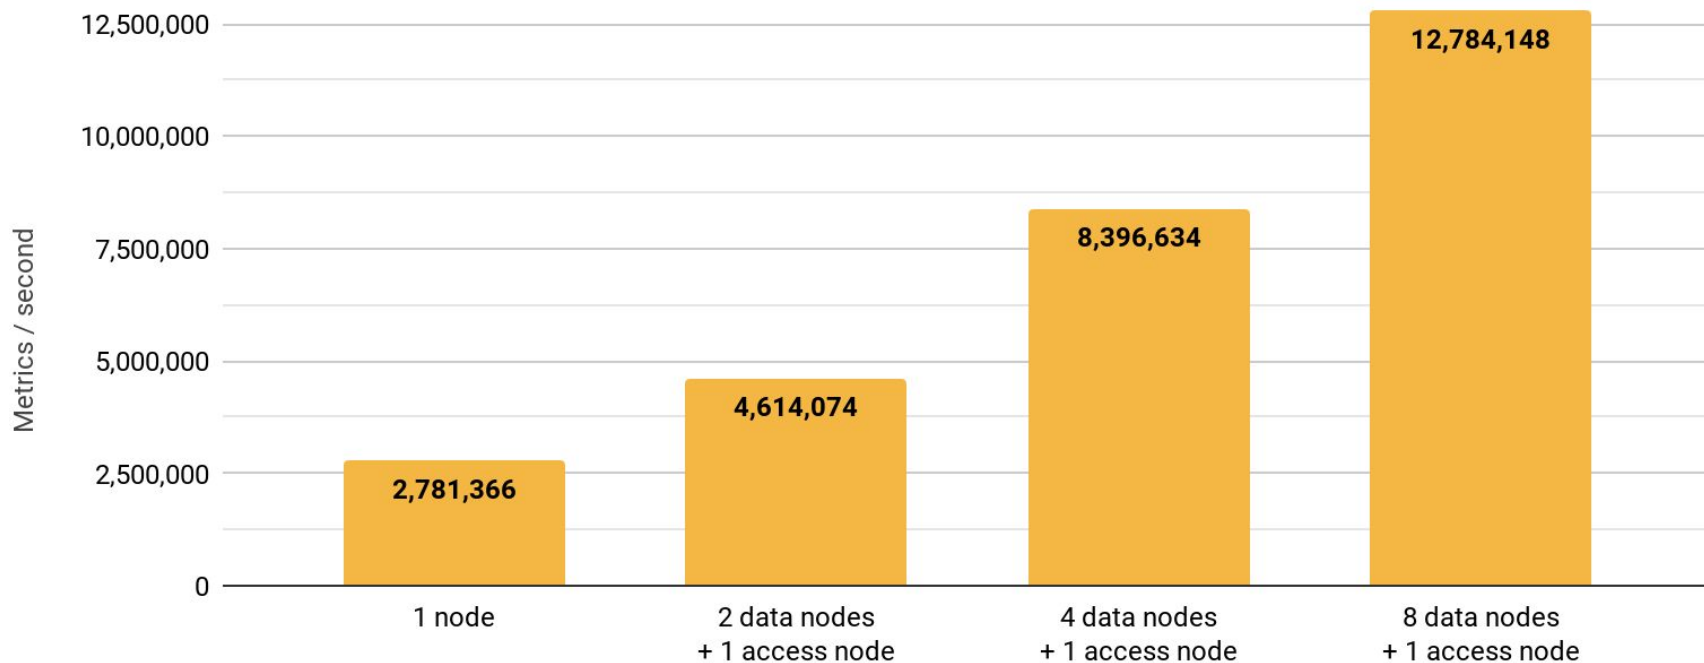
AN (Access Node) - узел доступа

DN (Data Node) - узел данных



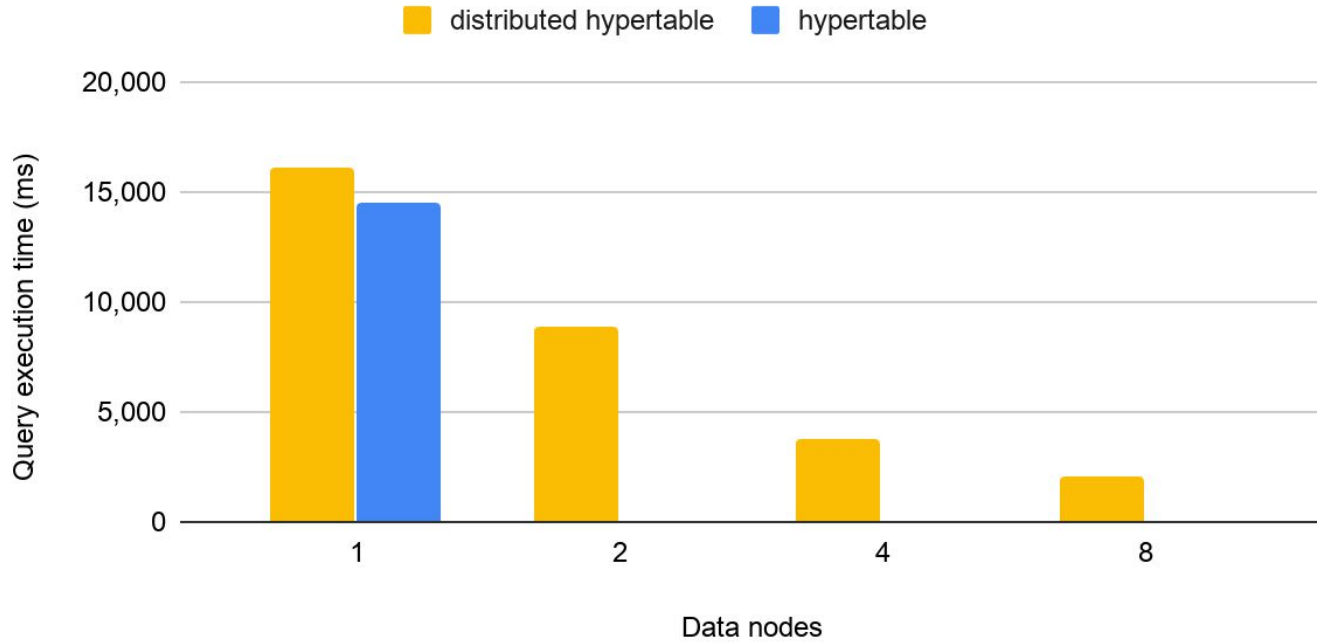
```
CREATE TABLE telemetries (  
    imei          TEXT          NOT NULL,  
    time         TIMESTAMPTZ   NOT NULL,  
    latitude     DOUBLE PRECISION NOT NULL,  
    longitude    DOUBLE PRECISION NOT NULL,  
    speed        SMALLINT      NOT NULL,  
    course       SMALLINT      NOT NULL,  
    CONSTRAINT telemetries_pkey PRIMARY KEY (imei, time)  
);  
  
SELECT * FROM add_data_node('data_node_1', host => 'pg_data_node_1');  
SELECT * FROM add_data_node('data_node_2', host => 'pg_data_node_2');  
SELECT * FROM add_data_node('data_node_3', host => 'pg_data_node_3');  
  
SELECT * FROM create_distributed_hypertable(  
    'telemetries', 'time', 'imei',  
    chunk_time_interval => INTERVAL '7 days'  
);  
  
INSERT INTO telemetries VALUES ...;  
  
SELECT  
    time_bucket('30 days', time) AS bucket,  
    imei,  
    avg(speed) AS avg, max(speed) AS max  
FROM telemetries  
WHERE speed > 0  
GROUP BY imei, bucket  
ORDER BY imei, bucket;
```

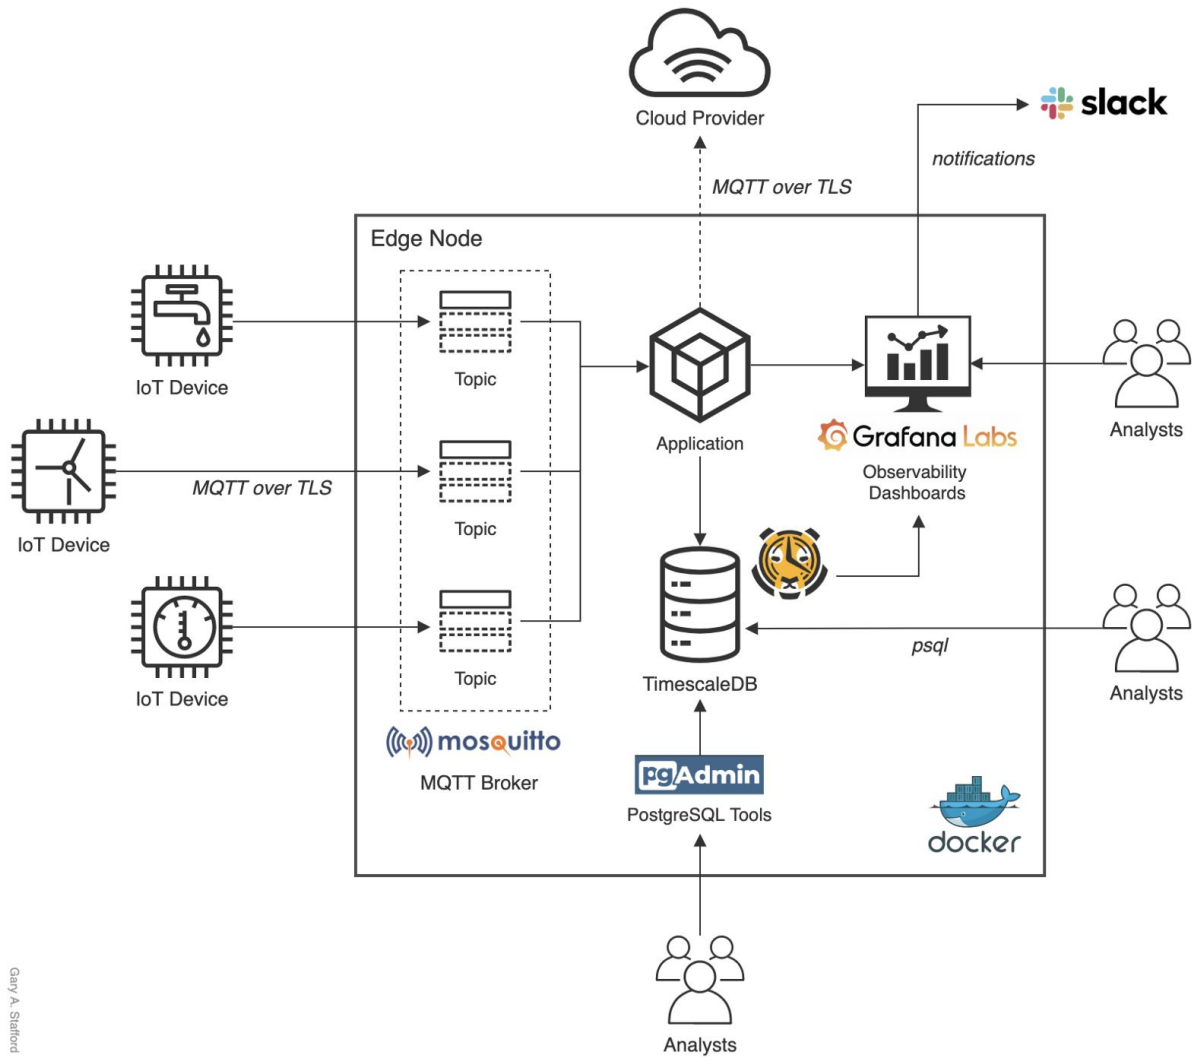
## Insert performance as cluster size increases





## Hourly average CPU usage per host over the past 24 hours (full aggregation)





# Limitations

Hypertable limitations:

- Foreign key constraints referencing a hypertable are not supported.
- Time dimensions (columns) used for partitioning cannot have NULL values.
- Unique indexes must include all columns that are partitioning dimensions.
- UPDATE statements that move values between partitions (chunks) are not supported.

Distributed hypertable limitations:

- Distributed scheduling of background jobs is not supported.
- Continuous aggregates are not supported.
- Joins on data nodes are not supported.
- Tables referenced by foreign key constraints in a distributed hypertable must be present on the access node and all data nodes.

<https://docs.timescale.com/timescaledb/latest/overview/limitations>

# ПРОФИТ?

- Настоящий SQL
- Экосистема PostgreSQL
- Масштабирование
- Открытый исходный код
- Свободное использование
- Огромное и дружелюбное комьюнити

Grafana ❤️ TimescaleDB

Zabbix ❤️ TimescaleDB

Ivan Muratov  
@binakot

Everything you need to know about [@TimescaleDB](#) compression - 13x on production database for vehicle telematics 🤓

```
waliot=# SELECT extversion
waliot=# FROM pg_extension
waliot=# WHERE extname = 'timescaledb';
-[ RECORD 1 ]-----
extversion | 2.4.2

waliot=# SELECT
waliot=# number_compressed_chunks,
waliot=# before_compression_total_bytes / after_compression_total_bytes as compression_ratio
waliot=# FROM hypertable_compression_stats('telemetrics');
-[ RECORD 1 ]-----
number_compressed_chunks | 246
compression_ratio        | 13

waliot=#
```

**Function pipelines: Building functional programming into PostgreSQL using custom operators**



# Полезняшки

@binakot   

<https://krd.dev>

<https://t.me/krddevdays>

<https://www.timescale.com>

<https://blog.timescale.com>

<https://youtu.be/vbJCq9PhSR0>

<https://youtu.be/1C2VGD90KGk>

<https://github.com/binakot/Multi-Node-TimescaleDB>

<https://hub.docker.com/r/binakot/postgresql-postgis-timescaledb>



**Спасибо за внимание!**





<https://github.com/binakot/Multi-Node-TimescaleDB>