



SQL Row Store vs Data Warehouse

Which Is Right for Your Application?

Robert Hodges
Percona Live Online 2020

Presenter Bio



Robert Hodges - Altinity CEO

30+ years on DBMS plus virtualization and security.

ClickHouse is DBMS #20



Altinity

www.altinity.com

The Enterprise Guide to ClickHouse

Major committer and community sponsor in US and Western Europe

What do we
mean by a row
store?



Key Features of MySQL

Single C++ binary

SQL language



Row storage

High concurrency

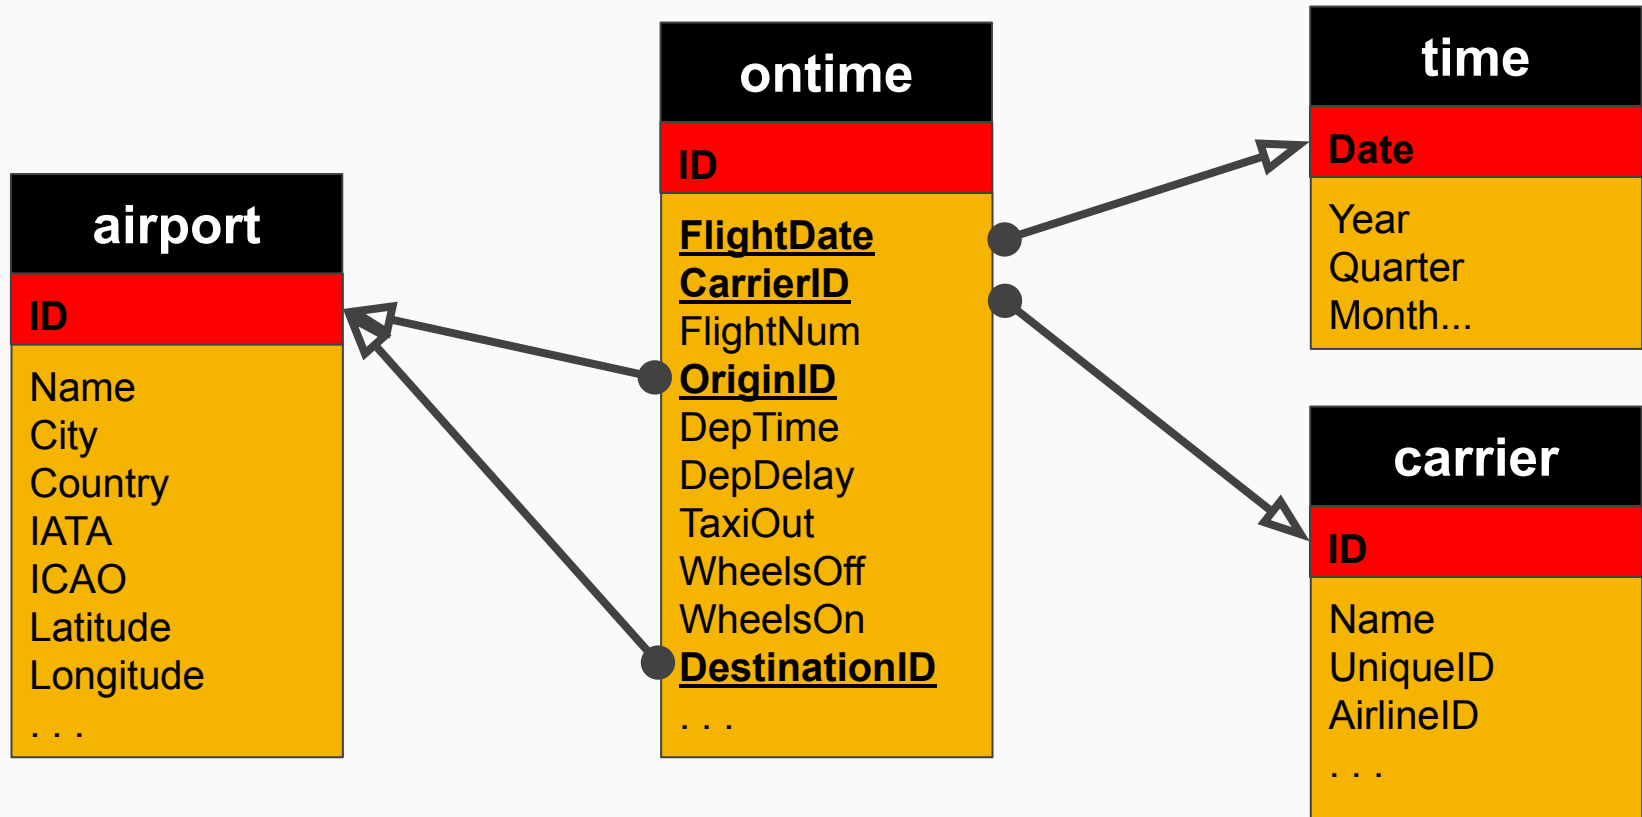
Single threaded query

Single server

GPL V2 license

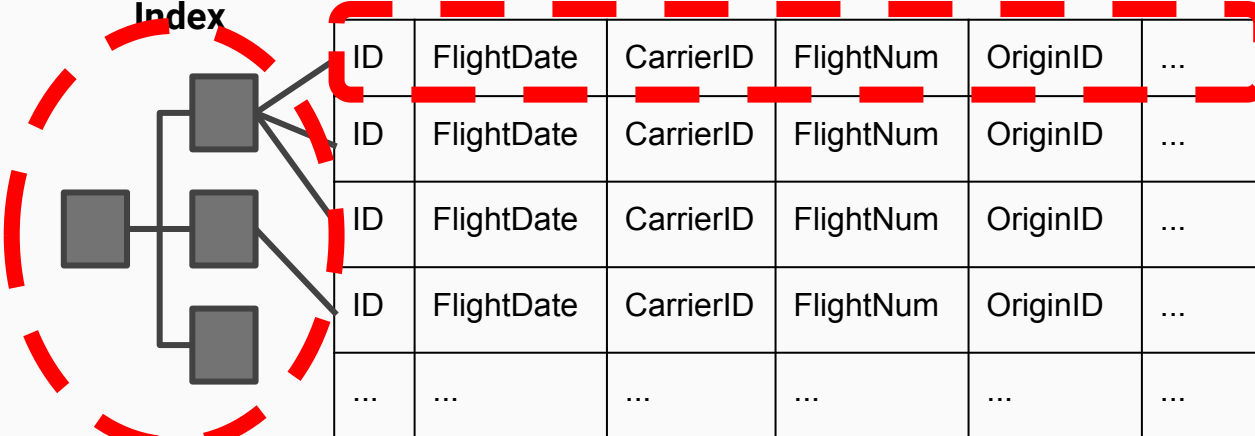


Airline flight management in MySQL



How row data is organized

Clustered Primary Key Index




ID	FlightDate	CarrierID	FlightNum	OriginID	...
ID	FlightDate	CarrierID	FlightNum	OriginID	...
ID	FlightDate	CarrierID	FlightNum	OriginID	...
ID	FlightDate	CarrierID	FlightNum	OriginID	...
...


No index? Read the entire table!

Secondary Indexes

FlightDate Index



CarrierID Index



...

Strengths and weaknesses of MySQL

OLTP

“Online Transaction Processing”

- (+) Lots of lookups
- (+) Lots of updates
- (+) High concurrency
- (+) Consistency critical

Analytic Queries

- (-) Very long tables
- (-) Very wide tables
- (-) Open ended questions
- (-) Lots of aggregates

MySQL servers can easily reach 1M QPS

How do data warehouses work?



Key features of ClickHouse

Single C++ binary

SQL language


Column storage

Compression and codecs

Vectorized query execution

MPP-enabled (shards/replicas)

Apache 2.0 license



Airline ontime table definition


```
CREATE TABLE airline.ontime (  
  `Year` UInt16, `Quarter` UInt8, `Month` UInt8, ...  
  `FlightDate` Date CODEC(DoubleDelta, ZSTD(1)),  
  `UniqueCarrier` LowCardinality(String) CODEC(ZSTD(1)),  
  `AirlineID` UInt32,  
  `Carrier` LowCardinality(String),  
  ... another 102 columns ...  
)  
ENGINE = MergeTree  
PARTITION BY FlightDate  
ORDER BY (Carrier, DepTime);
```

Table

Part


Part

Part



Rows match
PARTITION BY
expression

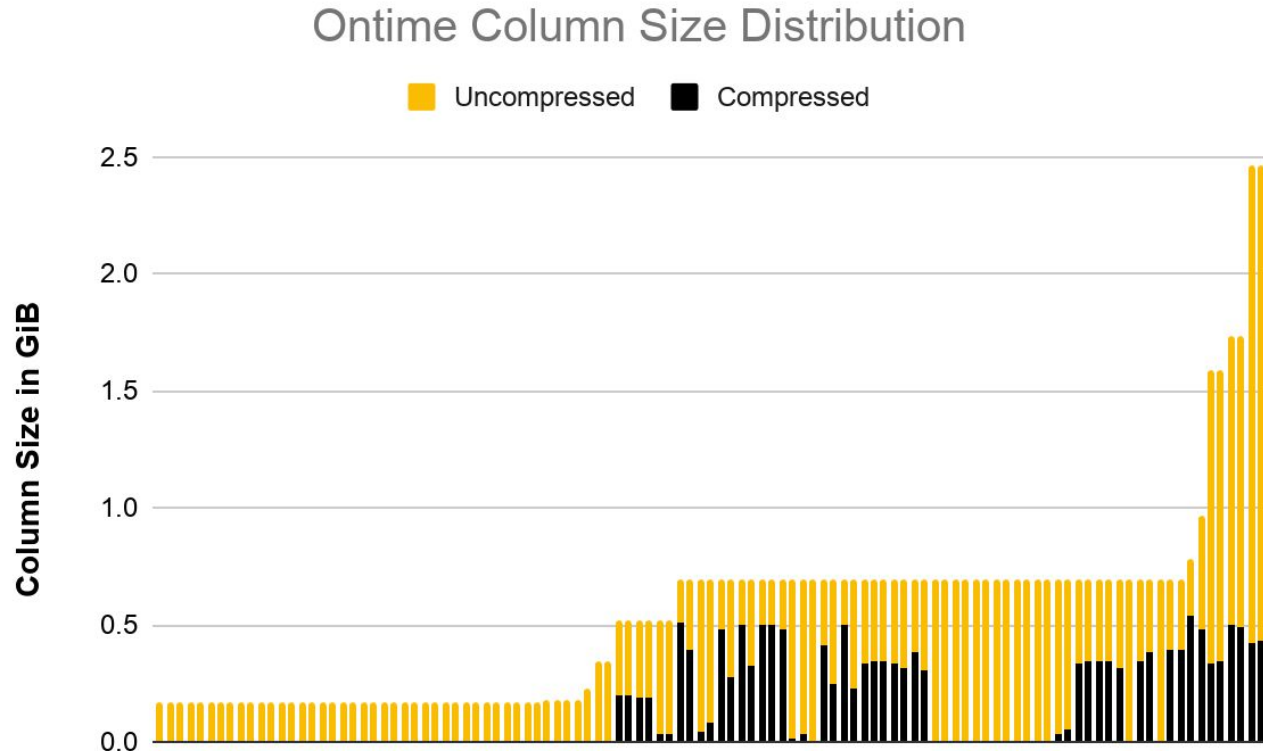
Columns sorted
on ORDER BY
columns



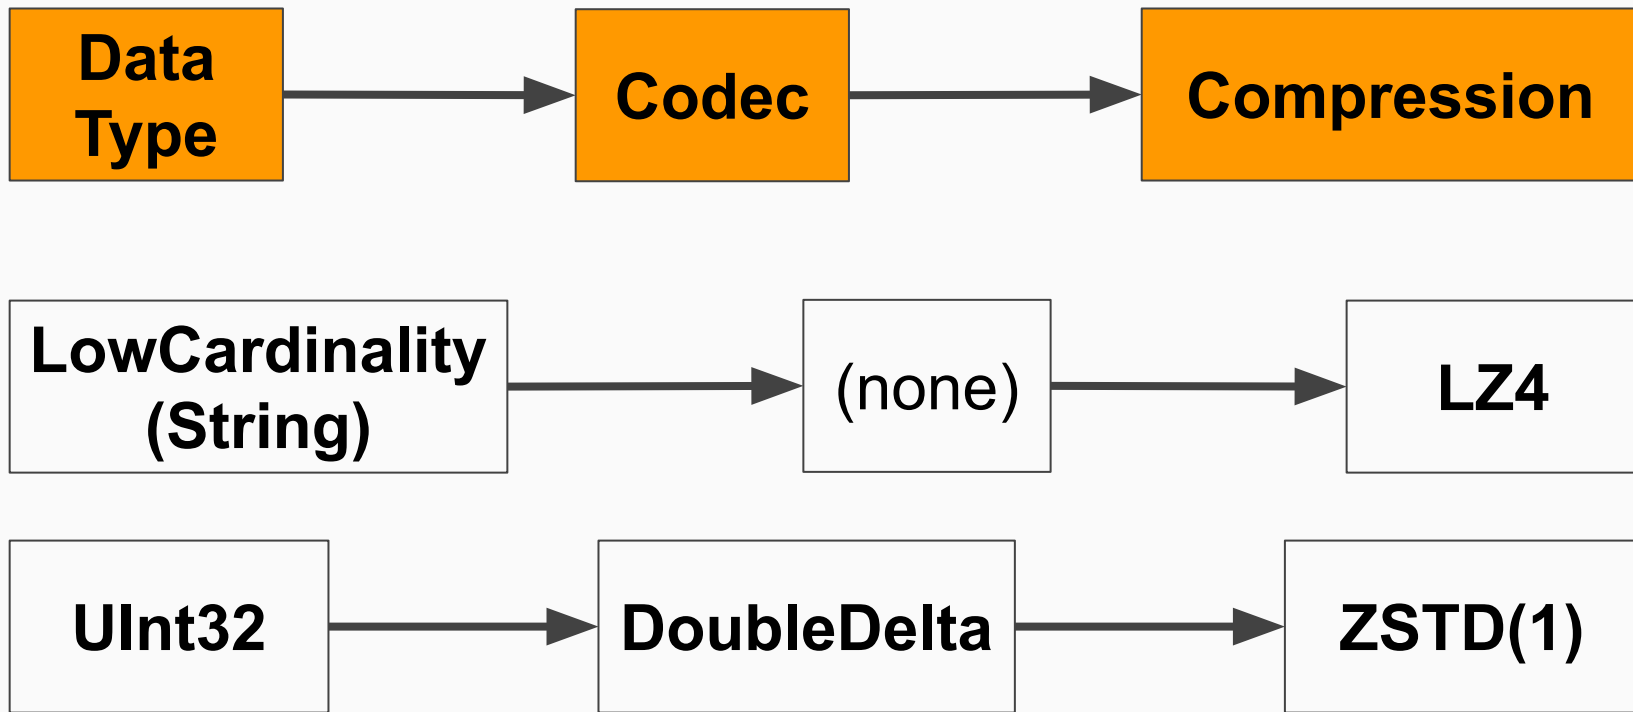
Compressed
block

Sparse index

There is no longer a penalty for wide tables

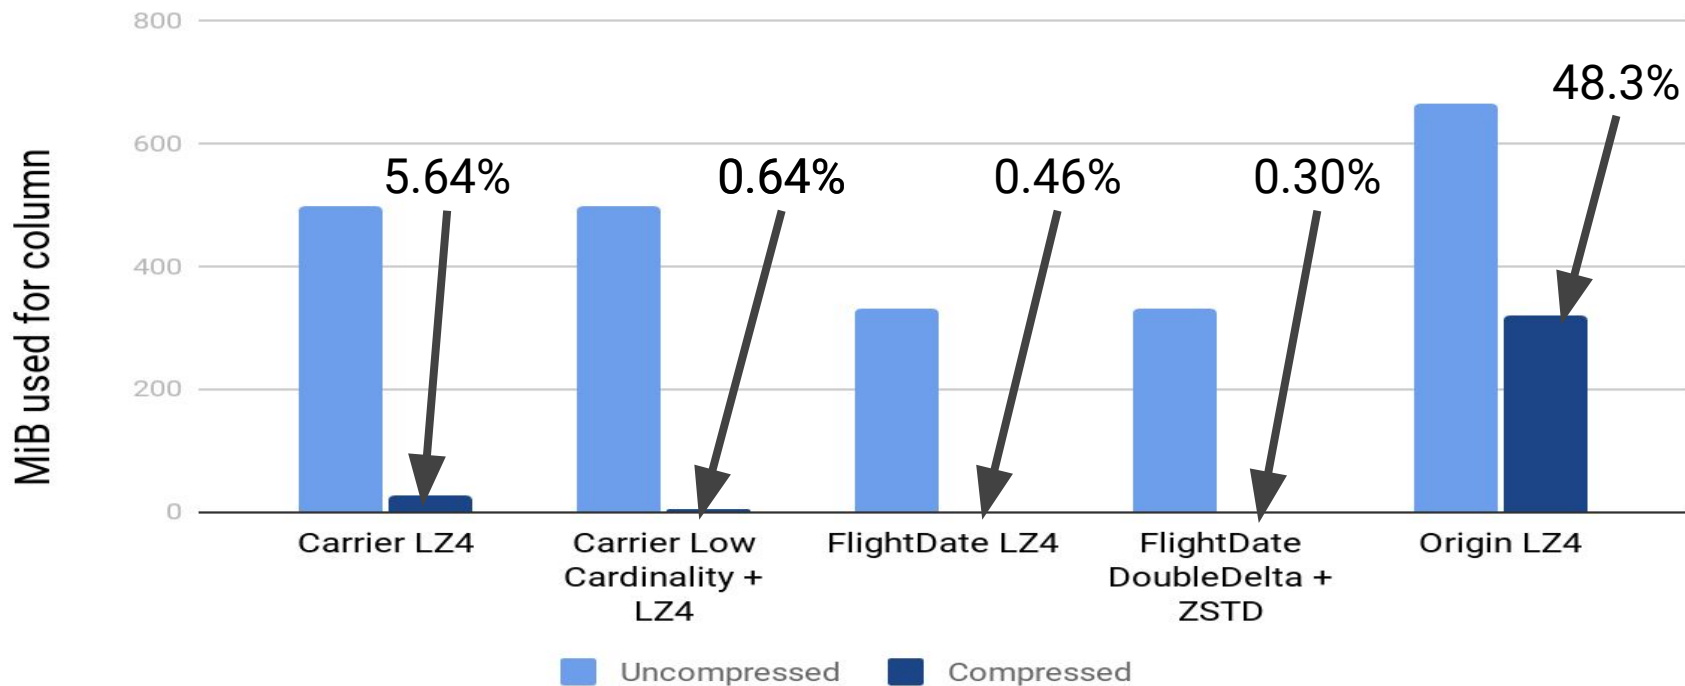


We make data even smaller by compressing



Data reductions from real data

Effects of Codecs and Compression



We can speed up queries using more CPUs

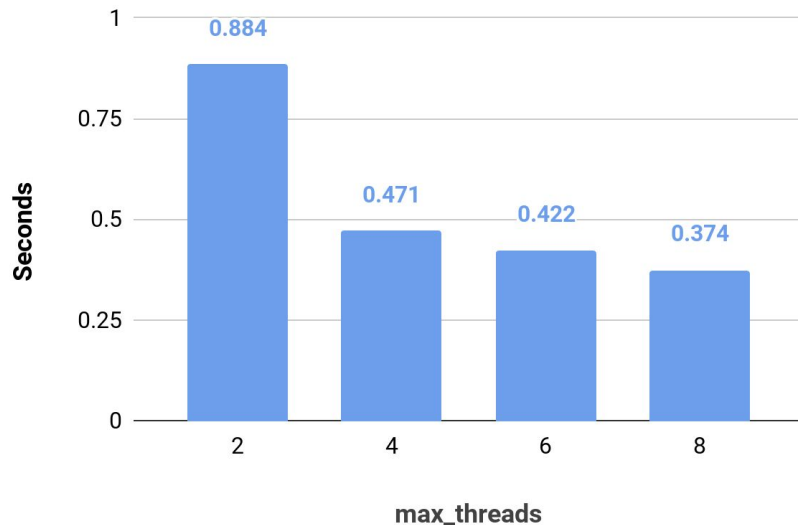
```
-- Find cancelled/delayed flights
SELECT toYear(FlightDate) year,
       sum(Cancelled)/count(*) cancelled,
       sum(DepDel15)/count(*) delayed_15
FROM airline.ontime
GROUP BY year ORDER BY year LIMIT 10

SET max_threads = 2

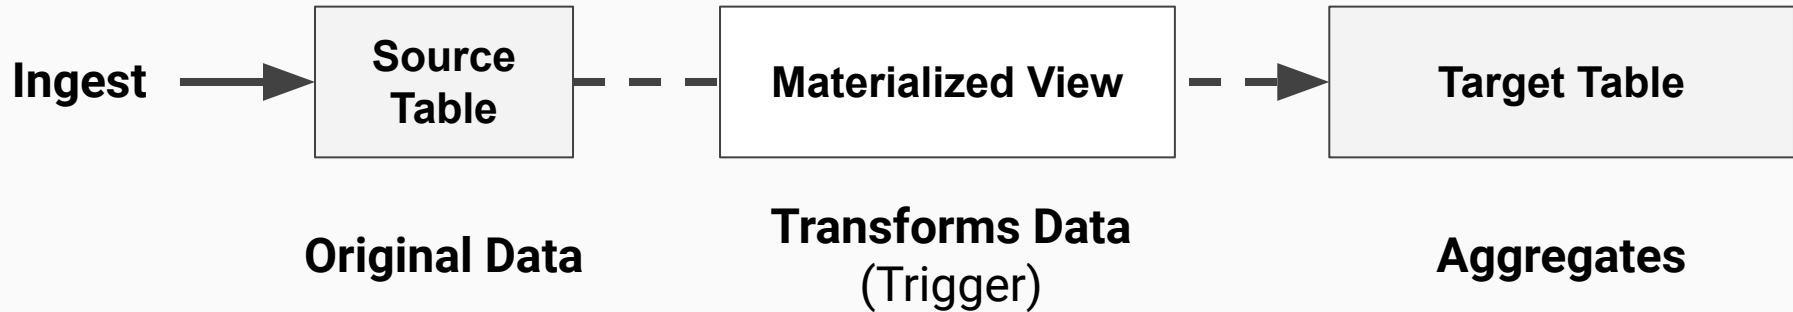
SET max_threads = 4

. . .
```


Query Response



Materialized views reduce data still more...



Materialized views restructure/reduce data



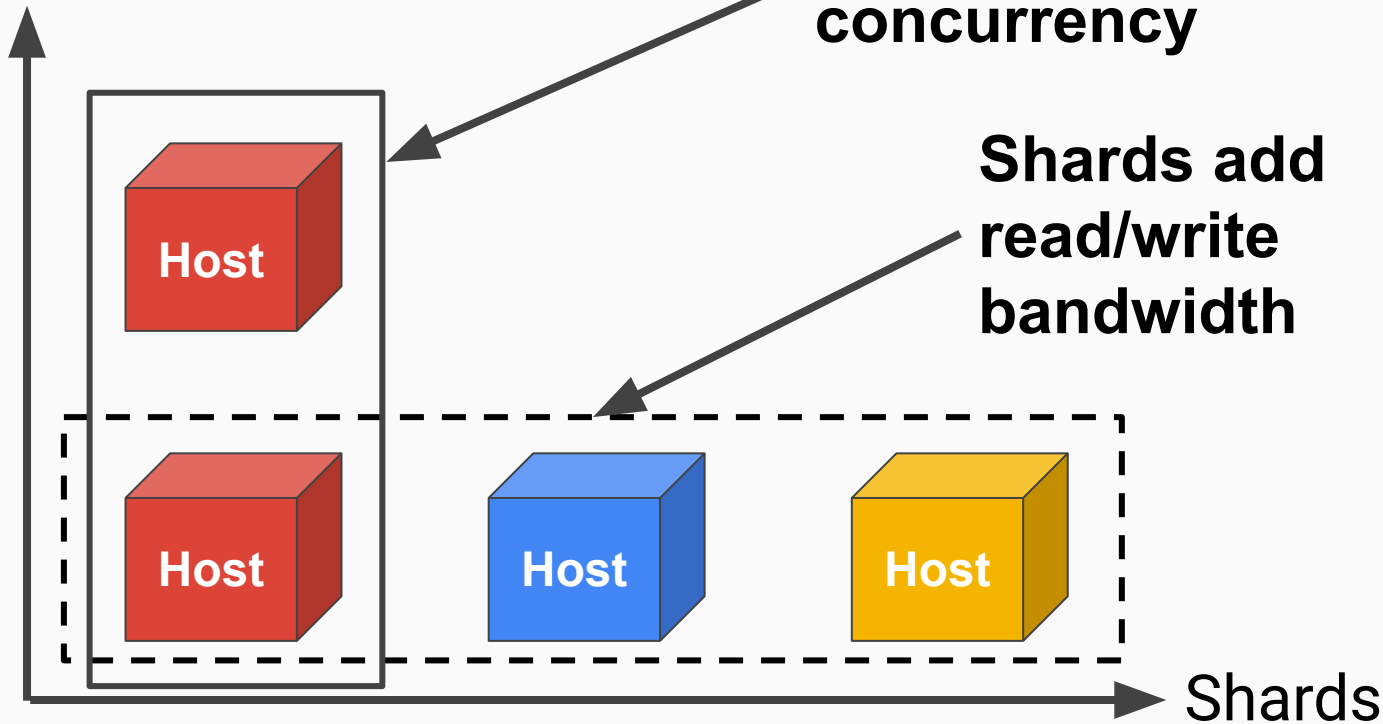
```
CREATE MATERIALIZED VIEW ontime_daily_cancelled_mv
ENGINE = SummingMergeTree
PARTITION BY tuple() ORDER BY (FlightDate, Carrier)
POPULATE
AS SELECT FlightDate, Carrier, count(*) AS flights,
        sum(Cancelled) AS cancelled,
        sum(DepDel15) AS delayed_15
FROM ontime GROUP BY FlightDate, Carrier
```

Still not fast enough? Add more servers!

Replicas

Replicas help with
concurrency

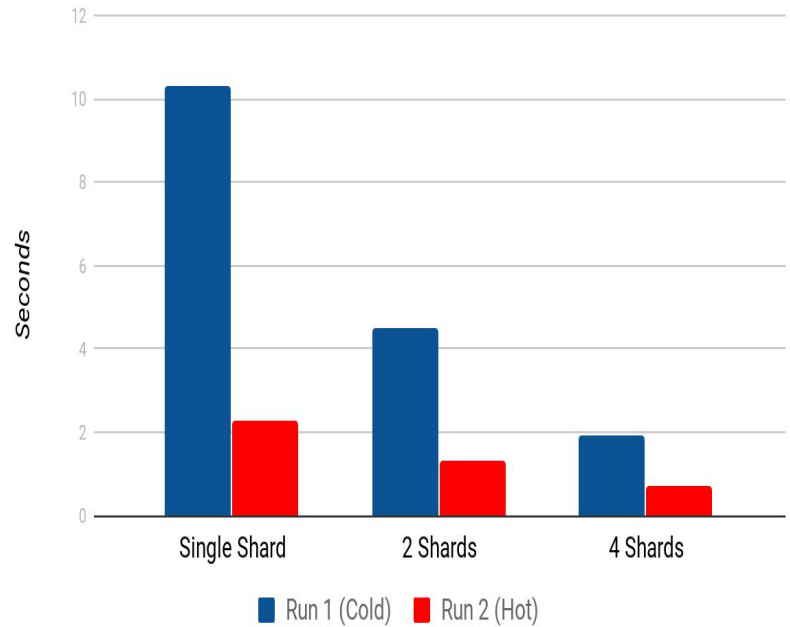
Shards add
read/write
bandwidth



ClickHouse distributes queries automatically

- Sharding and replication are built-in
- Replication is multi-master with eventual consistency
- Best case performance: linear with number of shards

Query over 173M rows



Strengths and weaknesses of ClickHouse

OLTP

(“Online Transaction Processing”)

- (-) Lots of lookups
- (-) Lots of updates
- (-) High concurrency
- (-) Consistency critical

Analytic Queries

- (+) Very long tables
- (+) Very wide tables
- (+) Open ended questions
- (+) Lots of aggregates

ClickHouse >> MySQL for analytic queries

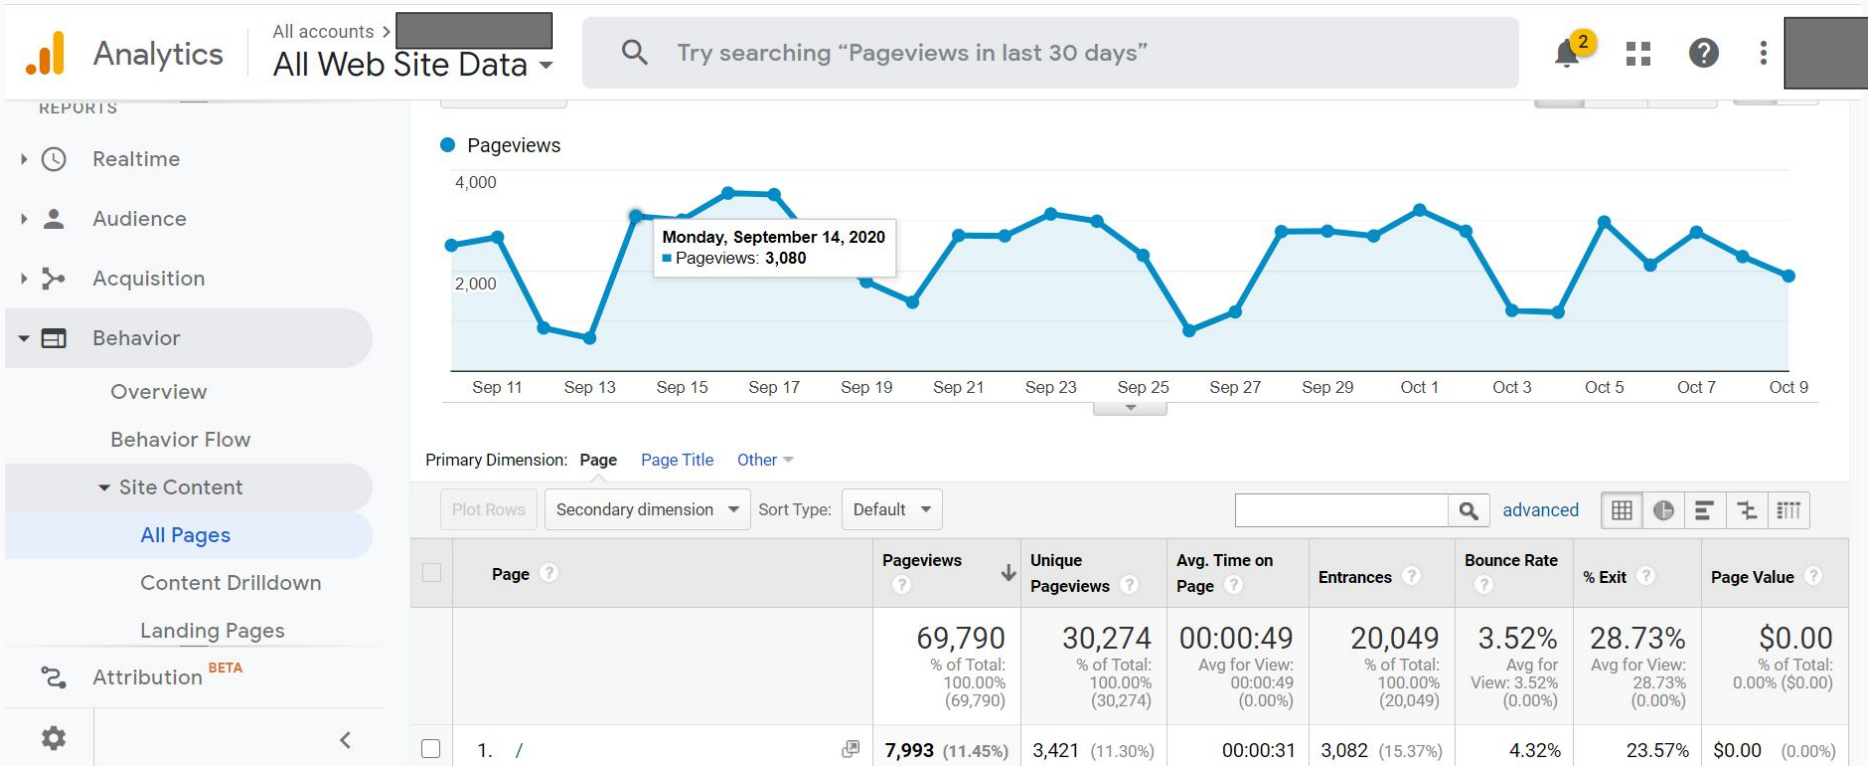
Wrap-up

Where MySQL shines -- transaction processing

The image shows a screenshot of the WordPress dashboard. At the top, there is a navigation bar with icons for Home, 1 New, Gallery, and Events. Below this is a sidebar menu with options: Dashboard, Home, Updates, Jetpack, Posts, Events, Media, NextGEN Gallery, Links, Pages, Comments (24), HubSpot, Feedback, News, Contact, and Portfolios. The main content area is titled "Dashboard" and contains several widgets:

- Message:** "The Page 'Events' uses the '/events' slug: the Events Calendar plugin will show its calendar in place of the page." with links to "Edit the Page slug" and "edit Events settings".
- Site Health Status:** "No information yet..." with a note that no Site Health information has been gathered yet and a link to "Visit the Site Health page".
- At a Glance:** A summary of site statistics: 213 Posts, 38 Pages, 113 Comments, and 24 Comments in moderation. It also notes "WordPress 5.5.1 running TheSaaSX-Child theme."
- Quick Draft:** A form for creating a new draft with fields for Title and Content (containing the text "What's on your mind?") and a "Save Draft" button.
- WordPress Events and News:** A section for upcoming events, featuring a discussion titled "Introduction to Contributing to" on Tuesday, Oct 13, 2020.

Where ClickHouse shines - analytic apps



A parting question...

Can one database handle both?

So far, **not really**.

Reason: Database and app complexity

Conclusions

- Row stores like MySQL are great for transaction processing
 - Lots of updates and small queries
 - Transactional integrity
- Column stores like ClickHouse are great for analytic queries
 - Wide, long tables
 - Open-ended questions
 - Lots of aggregated results

Thank you!

We're hiring

www.altinity.com

Email:

rhodges@altinity.com

ClickHouse:

<https://github.com/ClickHouse/ClickHouse>

MySQL:

<https://www.percona.com>