



# Streaming in a World of Legacy Applications

David Brimley <David@hazelcast.com>

Vladimir Schreiner <Vladimir@hazelcast.com>

# What is a Legacy System?



# No Documentation.



# The original devs are gone.






# Zero Tests



# Time






# THE MONEY MAKER!

\$ £ € ¥ ₨ Rs


Wayyyyyyyyyyyyyyy  
to scary to risk  
your career  
changing it.






Usually built  
around a core  
Relational  
Database.

# Typical Legacy Architecture







# Reasons for Change.

1.

Rooted in the present.

# Current Database.

A **current database** is a conventional **database** that stores data that is valid *now*. For example, if a user inserts "John Smith" into the Staff table of a current database, this asserts that the fact is valid *now* and until it is subsequently deleted. By contrast, a **temporal database** qualifies each row with a valid time *stamp*, valid time *period* or valid time *interval*. For example, we can assert the fact that "John Smith" was a member of staff during the period 1 June 2001 and *now*.

[https://en.wikipedia.org/wiki/Current\\_database](https://en.wikipedia.org/wiki/Current_database)

# Ignore Temporal Value of Data

# !Predictive Analysis

What if...

- Time Period Definitions
  - Start & End Times
- Temporal Primary Keys
- Temporal Referential Integrity
- Temporal Predicates
  - Contains
  - Overlaps
  - Equals
  - Precedes
  - Succeeds
  - Immediately Precedes
  - Immediately Succeeds
- Time Slice Queries

# SQL:2011

# 2.

## Answer Questions Now.

# Store then Process



# Process then Store

# 3.

## Distributing Events

Most RDBMS are  
passive stores

# Push over Pull

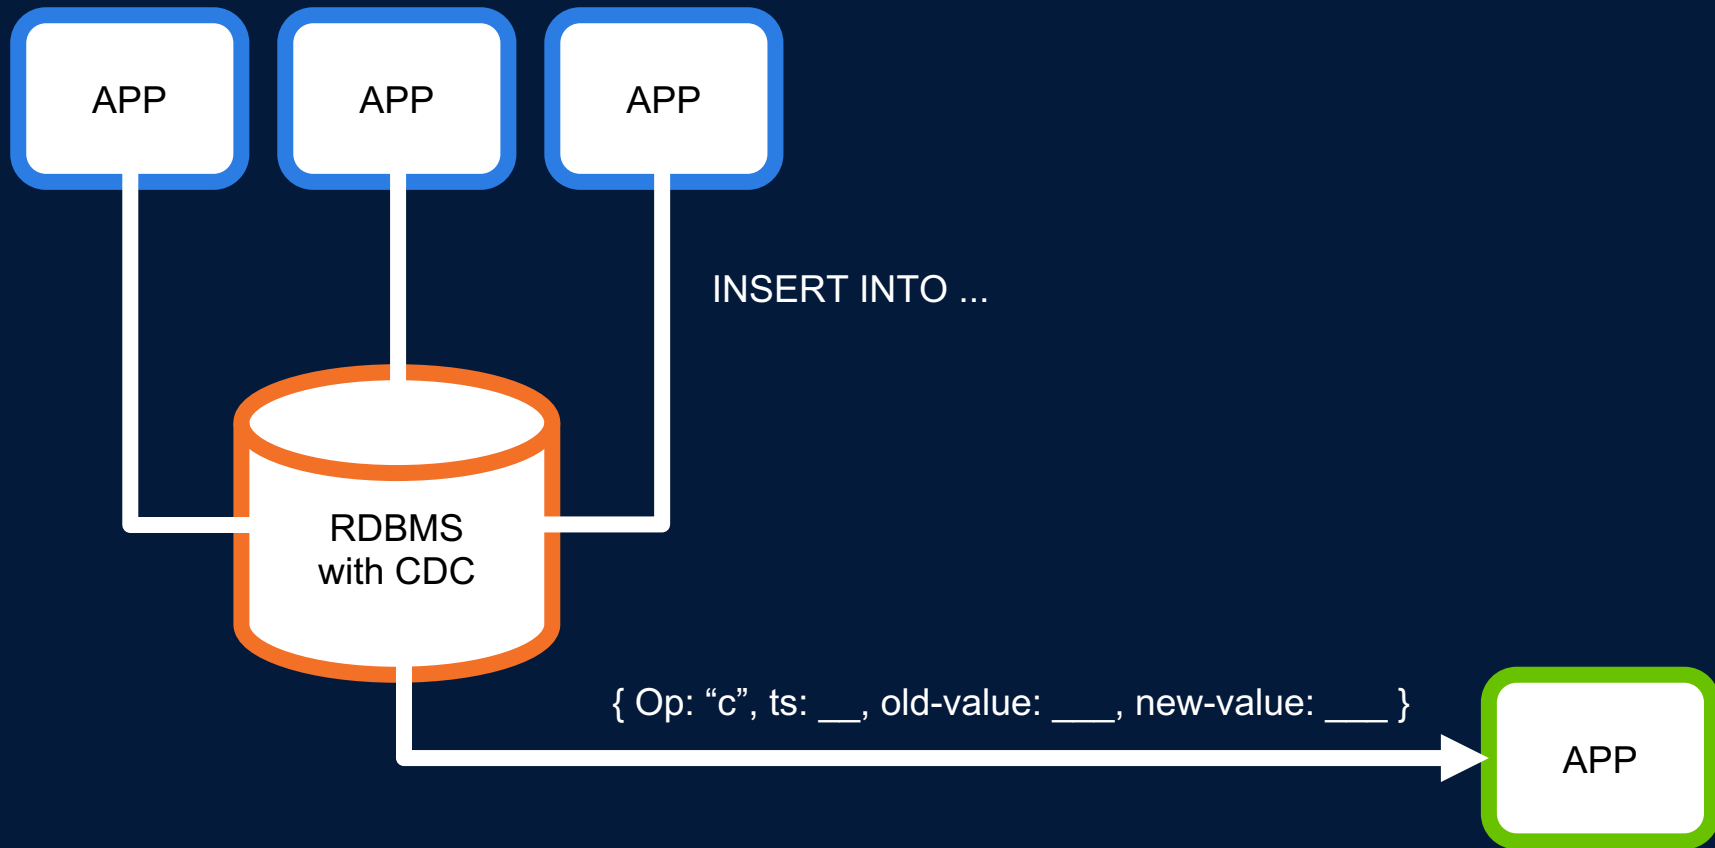
Easier to  
integrate new  
systems




# IMPEDANCE MISMATCH

Change Data Capture  
(CDC) turns legacy  
database to a source of  
event stream.











# Travelling Back in Time

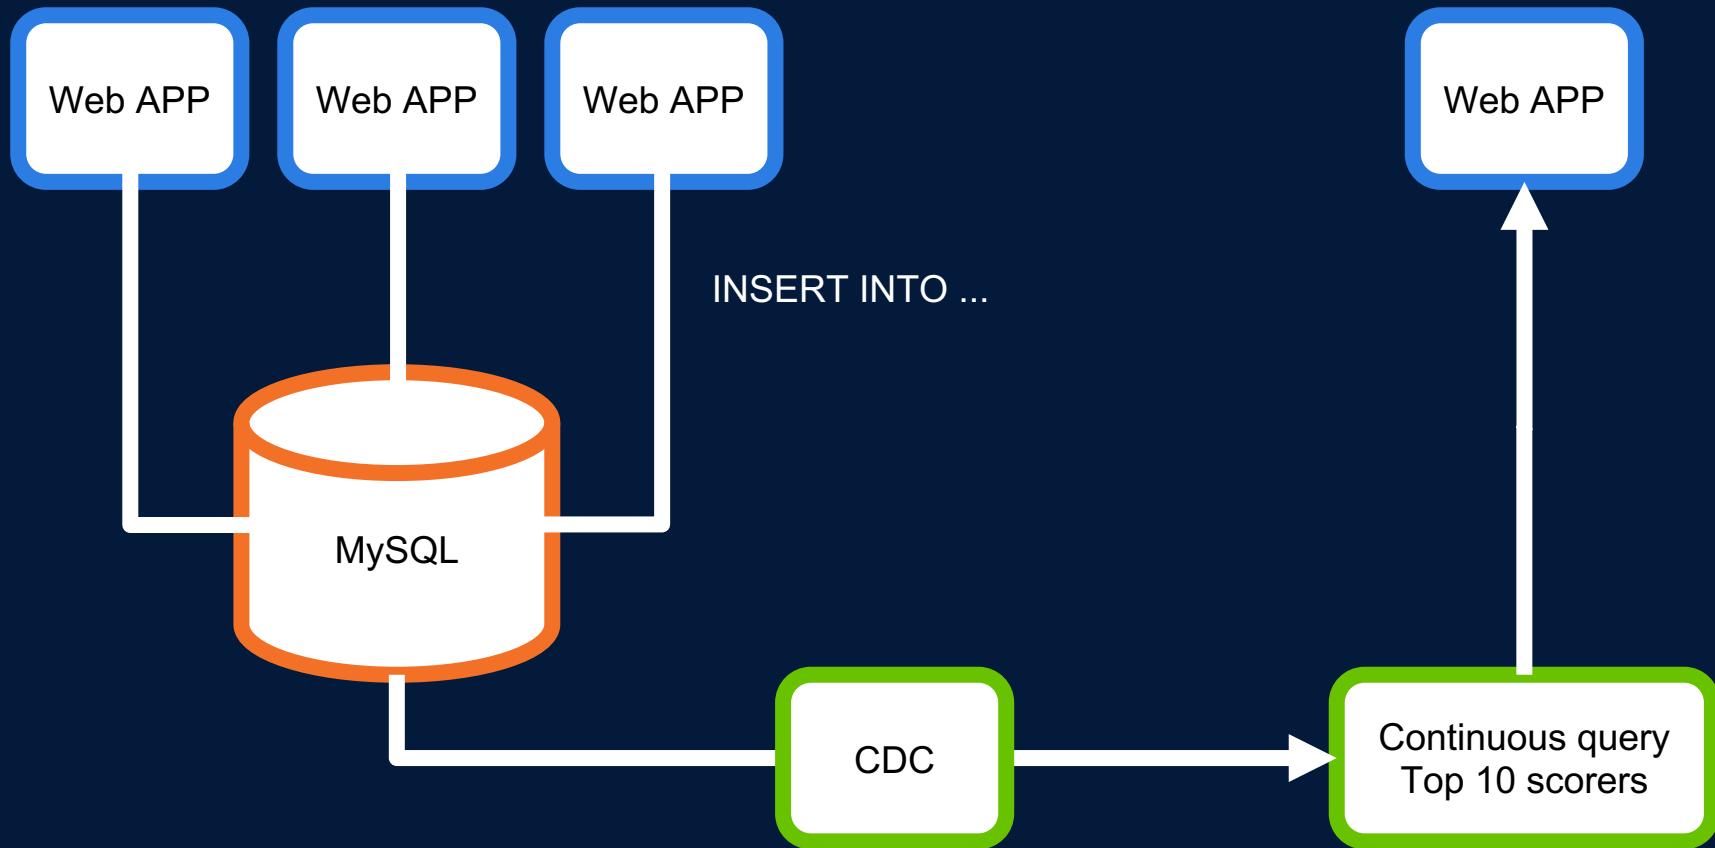
<https://github.com/hazelcast/hazelcast-jet-demos/tree/master/debezium-cdc-without-kafka>



LAMP stack  
GBs of data  
Mostly OLTP,  
OLAP after  
hours



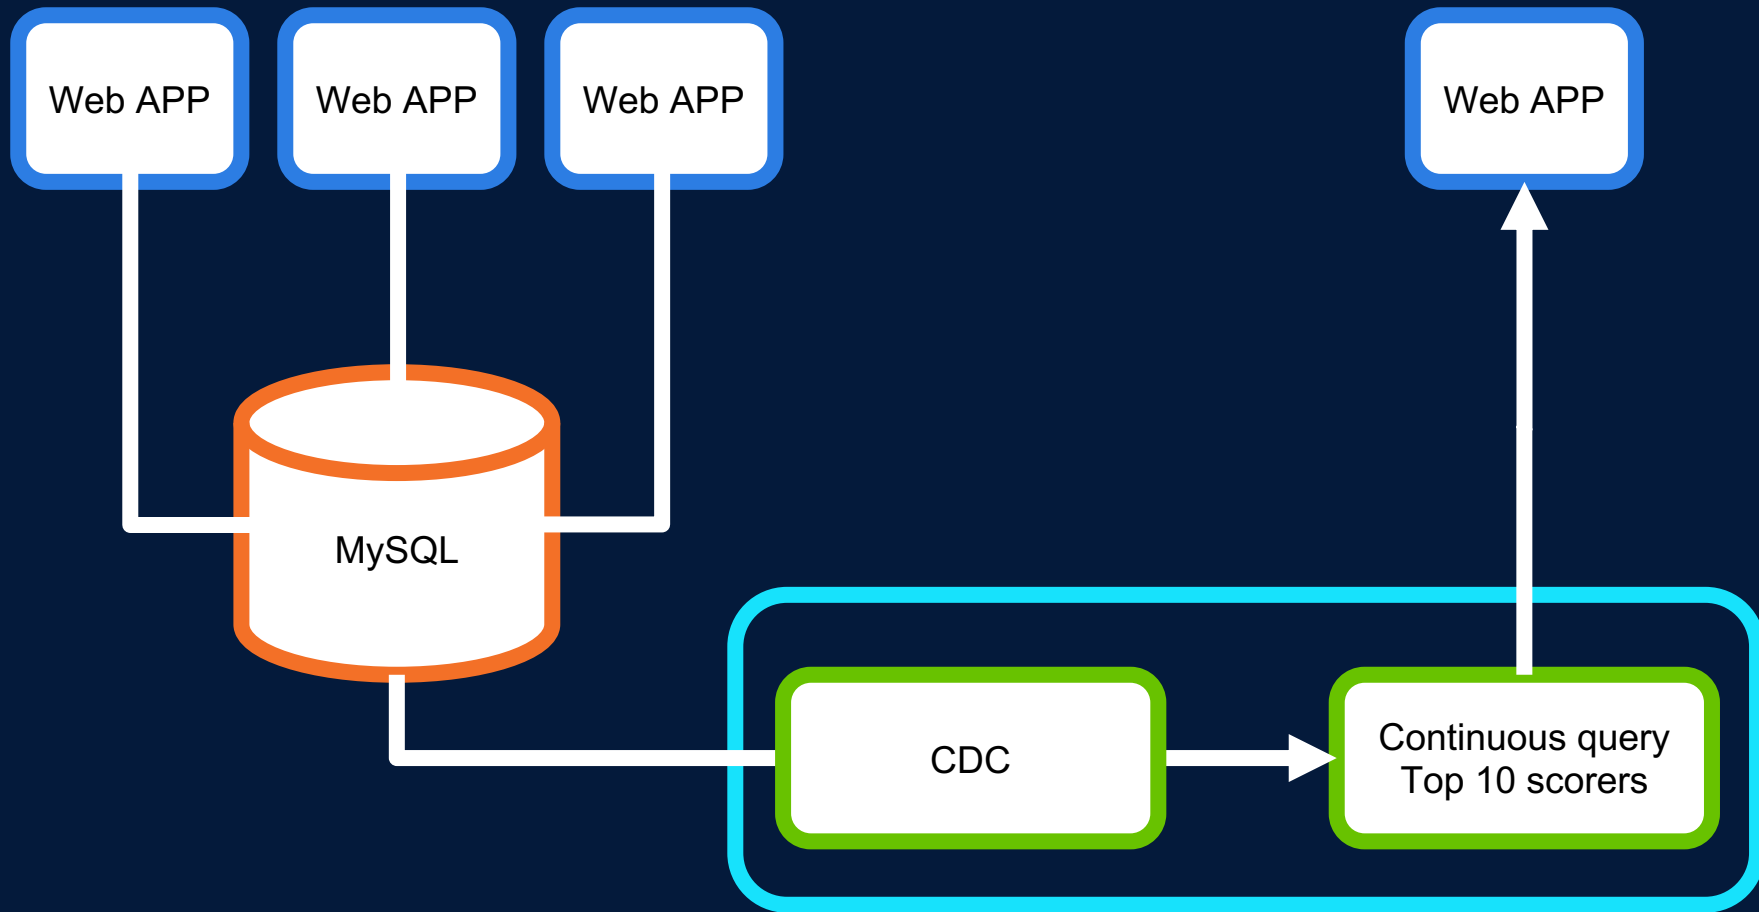
Real-time  
updates =  
Real-time  
expectations



Open-source library with  
stream processor,  
connectors including CDC  
and a in-memory key-  
value store.

Single Java Binary  
Elastic Clustering  
No ZooKeeper, HDFS..  
Just Java 8 and above

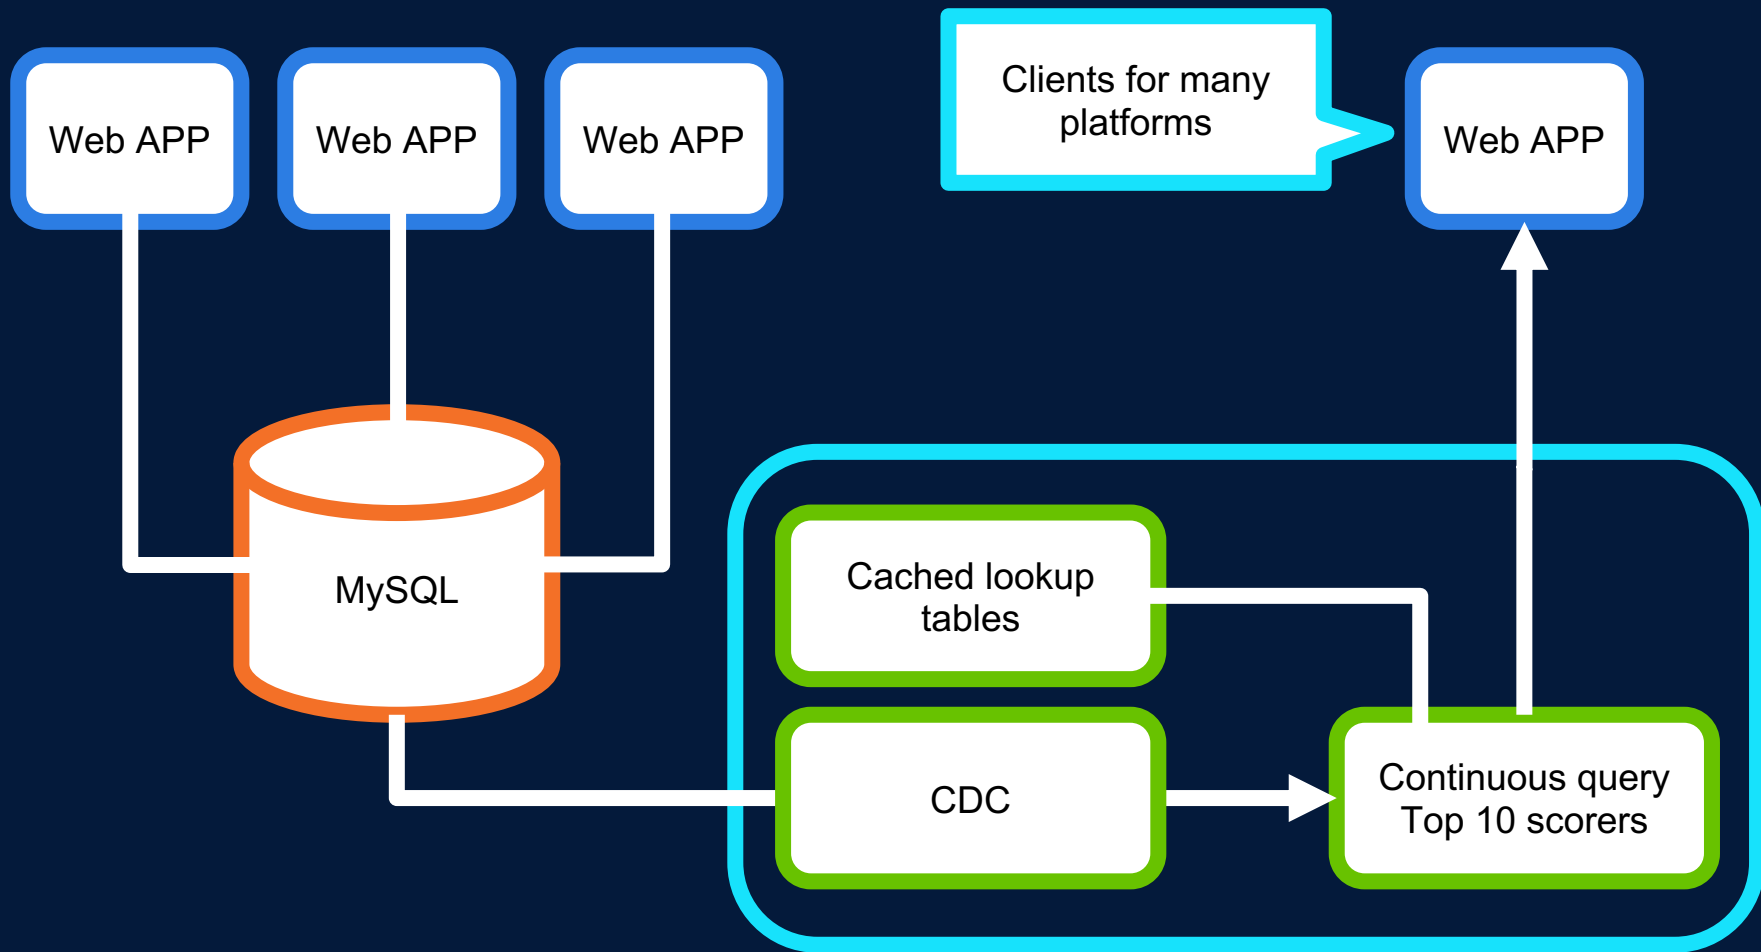




# Hazelcast Jet runs Debezium for CDC

```
# The following can be used as easy to replay  
# backup logs or for replication.  
server-id                = 1  
log_bin                  = /var/log/mysql/mysql-bin.log  
binlog_format           = row  
binlog_row_image        = full  
expire_logs_days       = 2
```

Relational databases  
usually keep shorter  
history, compared to  
dedicated log-based  
storages.



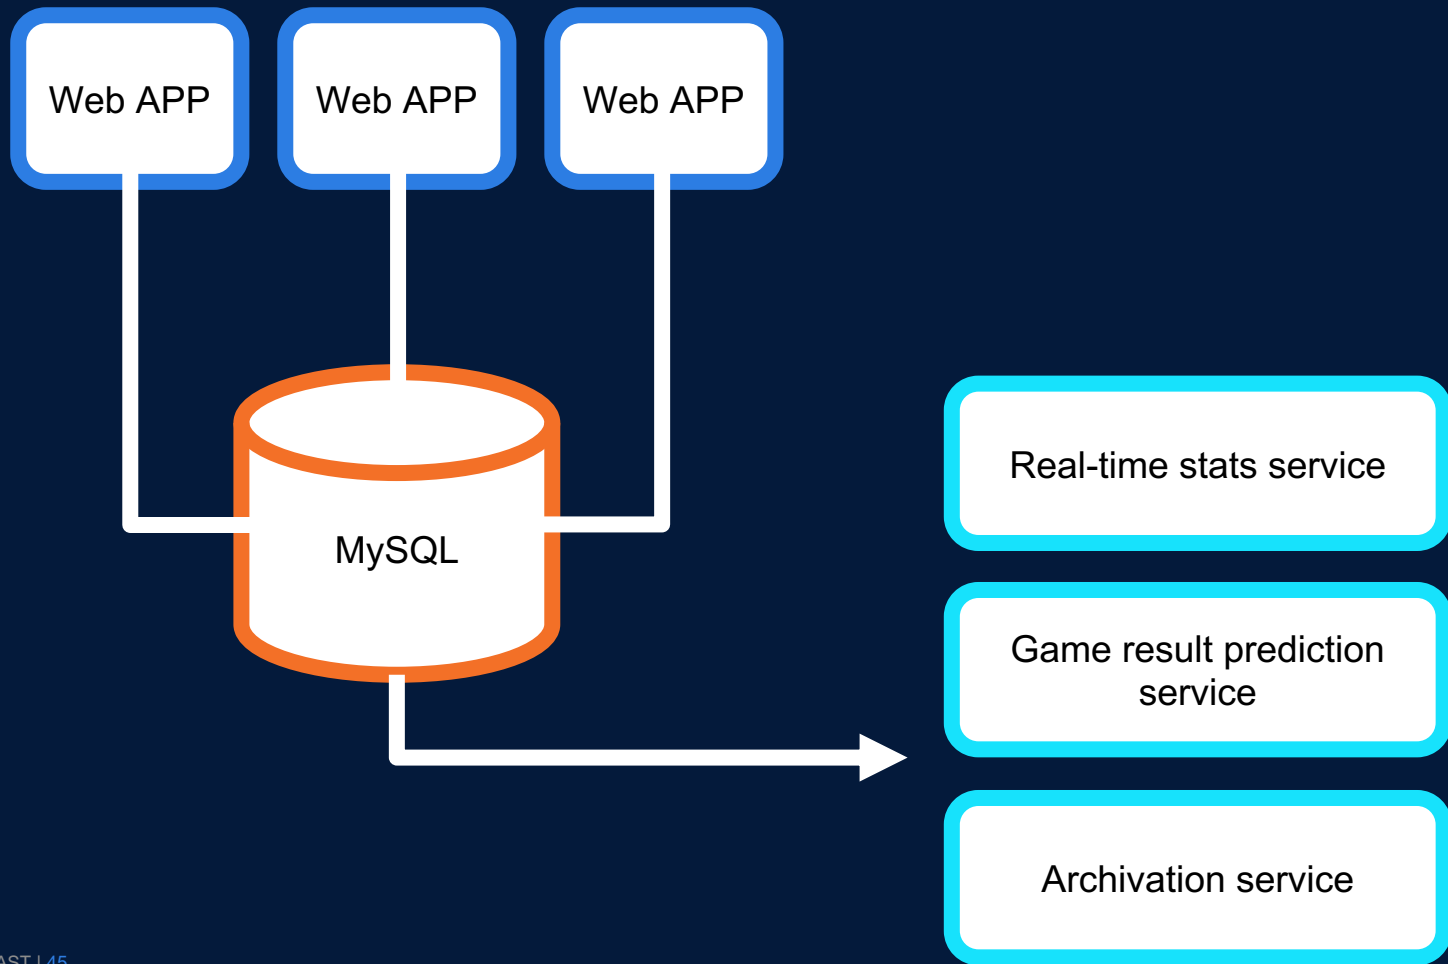
# Event-driven end2end

CDC + SPE + CACHE =

Materialized Views  
Offloaded from the DB

# Modularization Microservices





# Scaling

# Conclusions

Extract change stream from your legacy stack using CDC.

Use CDC and stream processing to create materialized views outside the database to reduce the database load and modularize your architecture.

Hazelcast Jet is an open-source Java library to do the job.

CDC

Streaming

Caching

# Questions

Visit our stand on 3rd floor.

[jet.hazelcast.org](http://jet.hazelcast.org)