




# Stream Processing in Java

Presented by:

Vladimír Schreiner <[vladimir@hazelcast.com](mailto:vladimir@hazelcast.com)>  
[@voloda](https://twitter.com/voloda)


## > Housekeeping




# > Agenda

- Modern DBs x Streaming
- Concern #1 - Time
- Concern #2 - Connectors
- Concern #3 – Scaling?


# > System Health Monitoring



10:22:01 12ms  
10:22:01 15ms  
10:22:02 10ms  
...



10:22:00 150ms  
10:22:02 159ms  
10:22:02 170ms  
...



2

```
A: SELECT AVG(responseTime)
FROM logs
WHERE timestamp
BETWEEN NOW()
AND NOW() - INTERVAL 1 DAY;
```

```
B: ...
AND NOW() - INTERVAL 1 SECOND;
```

3  $B / A > 1,1$  -> Alert event

## > DB tuned for append-only tables


- Stream: ordered sequence of immutable records
  - Log records, clicks, IoT readings, business events, ...
- Stream Processing: querying the stream continuously
  - Continuously transform, join, aggregate, ...

# ➤ Focal areas of Streaming


1. Time as a first class citizen
2. Connectivity
3. Scale

# > Time


# > When should we get the results?



10:22:01 12ms  
10:22:01 15ms  
10:22:02 10ms  
...



10:22:00 150ms  
10:22:02 159ms  
10:22:02 170ms  
...



2

```
A: SELECT AVG(responseTime)
FROM logs
WHERE timestamp
BETWEEN NOW()
AND NOW() - INTERVAL 1 DAY;
```

```
B: ...
AND NOW() - INTERVAL 1 SECOND;
```

3

$B / A > 1,1$  -> Alert event




## > Stream Processing is “push”

- Database allows clients to **pull** data by querying it's state
- Streaming Engine runs a continuous query and **pushes** updates to consumers
- Continuous programming model decouples query submission from result materialization!
  - Computations driven by input data, not by query submission
  - Low latency and correct results!

## > Windows to define data ranges

Windows define the ranges in the append only table

Engine runs the query as soon as it has complete data for the window



## > New API to define windows

```
SELECT AVG(resoponseTime)
FROM Logs
GROUP BY SLIDING(timestamp,
INTERVAL '1' SECOND, INTERVAL '1' DAY)
```

Results are published after the engine recieved and processed all input data for the **time window**.

## > When the window completes?

App 1

App 2

Stream Processor

10:22:00 150ms

10:22:01 159ms

10:22:02 170ms

*Ok, I can publish results  
for 10:22:00 now..*

10:22:00 12ms

10:22:00 15ms

10:22:02 10ms

*WTF???*

Time (as observed by stream processor)

## > Strategies for unordered data

- Waiting for stragglers

*"Maybe somebody is late, let's waiting another \_\_\_ before publishing results"*

- Publish early results

*"Based on the data seen so far, the result is \_\_\_"*

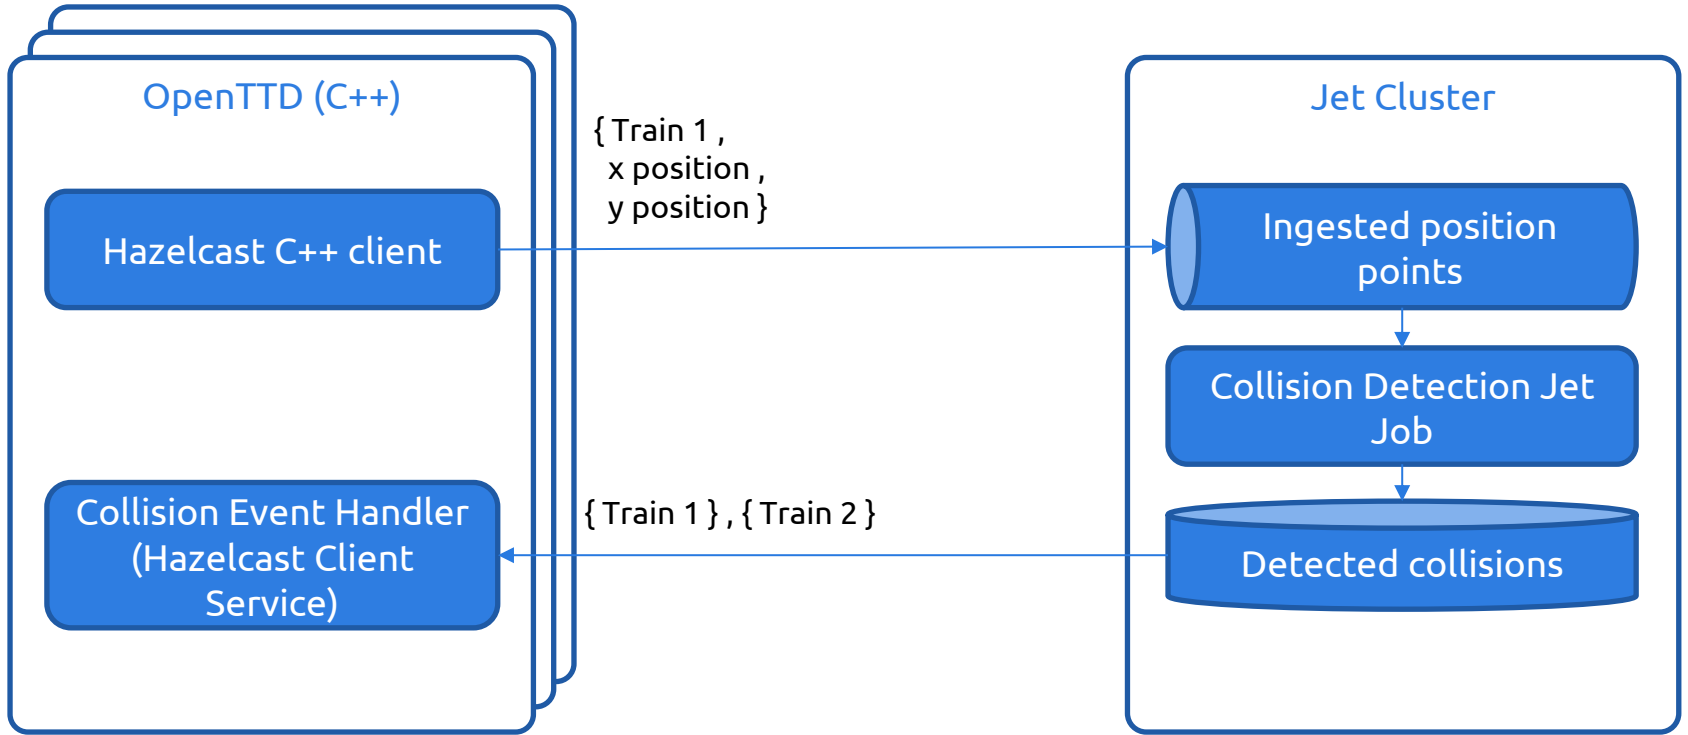
## > Time - Summary

- Streaming queries run continuously and push results to subscribers.
  - **Event-driven querying for lower latency**
- Declarative API to for data driven "triggers"
  - Windowing - which data are required for the computation
  - When to publish the results


## > Use Case: Analytics and Decision Making

- Real-time dashboards
- Stats (gaming, infrastructure monitoring)
- Decision making
- Recommendations
- Prediction - often based on algorithmic prediction (push stream through ML model)
- Complex Event Processing
- Moving average

# > Train Demo!








# > Connectors

# Typical Legacy Architecture























# > One needs stream of data for streaming!

- Stream is a sequence of immutable events, the append-only table
- To enable streaming, we need applications that:
  - Produces events
  - Deliver events to a streaming engine
  - E.g.: using client (agent), publish changes to a message broker

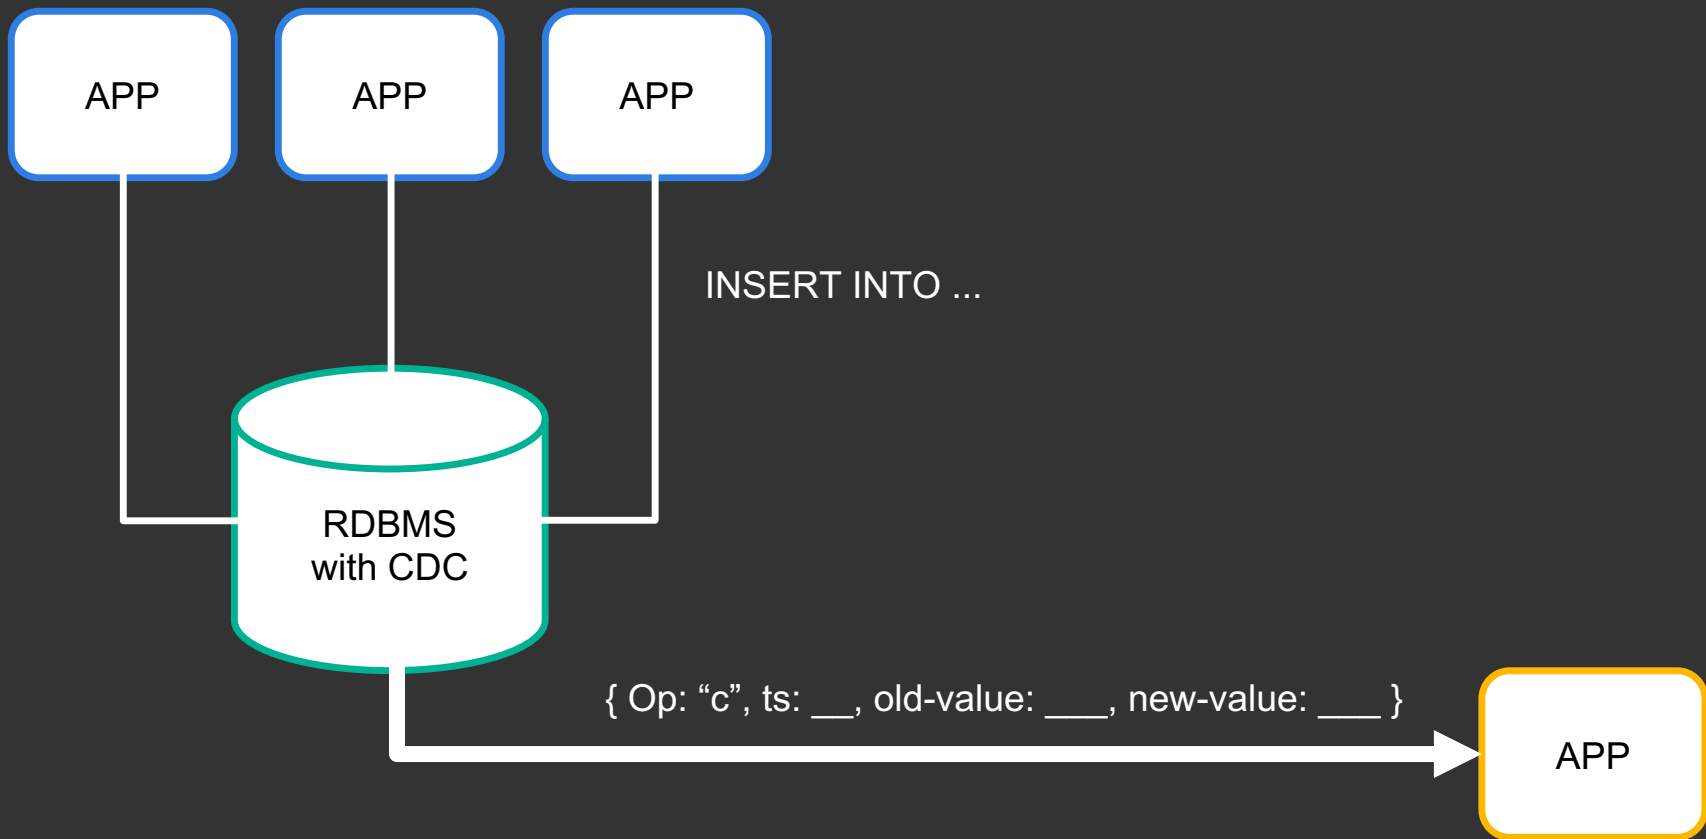
# > In the meanwhile in a real world

DB Systems in Ranking, January 2020

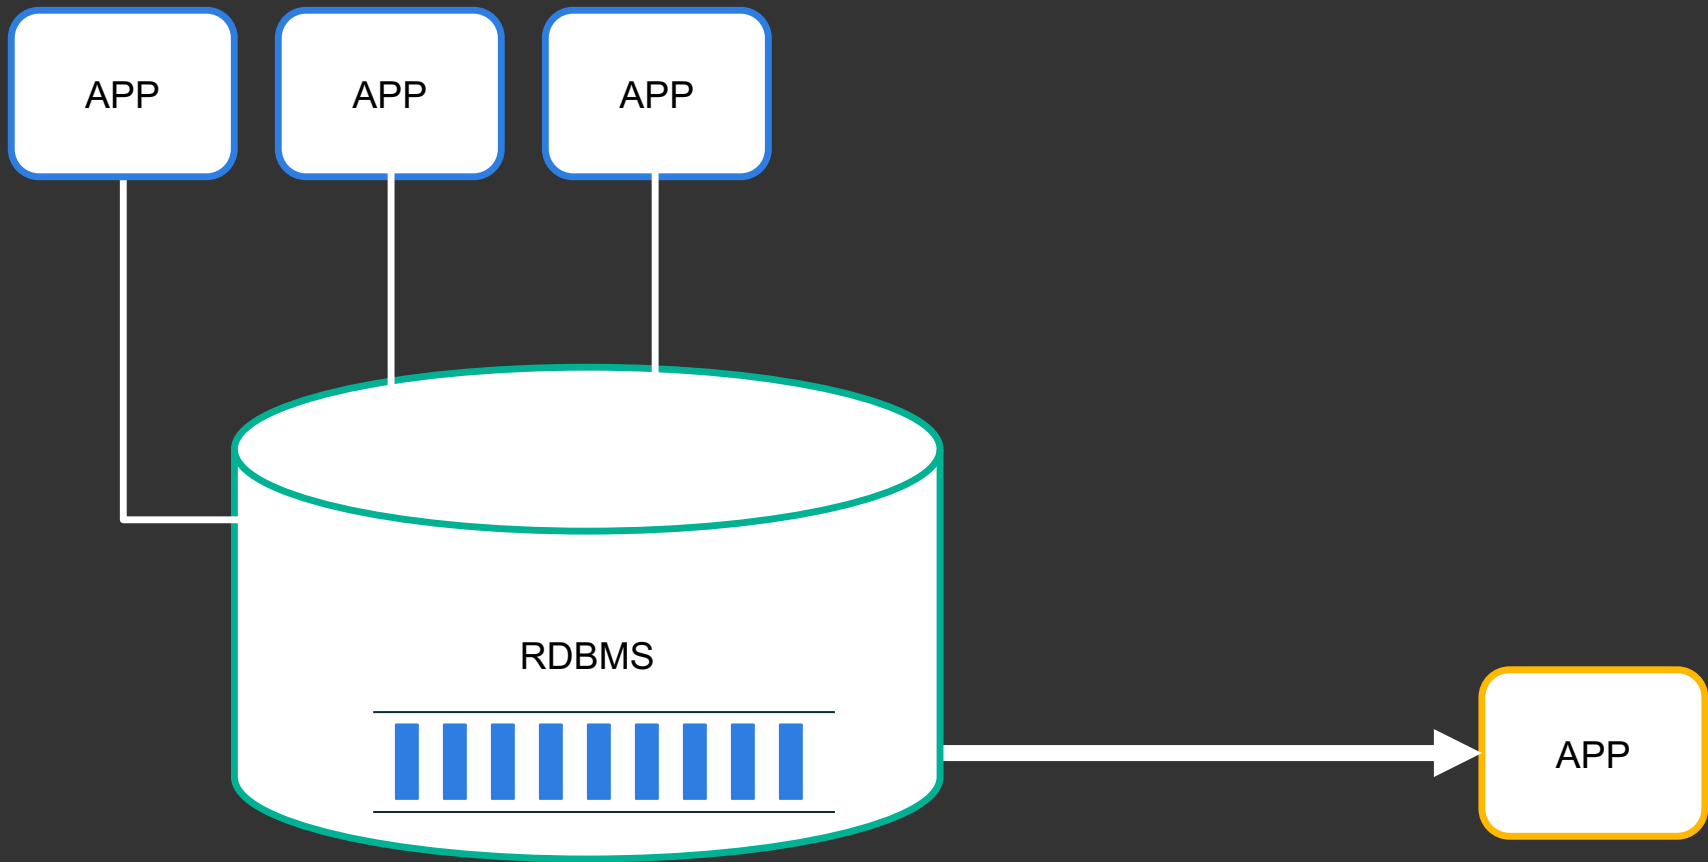
| Rank     |   |  | DBMS   | Database Model   | Score    |          |          |
|----------|---|--|--|--|----------|----------|----------|
| Jan 2020 | Dec 2019  | Jan 2019   |  |  | Jan 2020 | Dec 2019 | Jan 2019 |
| 1.       | 1.  | 1.   | Oracle                | Relational, Multi-model     | 1346.68  | +0.29    | +77.85   |
| 2.       | 2.  | 2.   | MySQL                 | Relational, Multi-model     | 1274.65  | -1.01    | +120.39  |
| 3.       | 3.  | 3.   | Microsoft SQL Server  | Relational, Multi-model     | 1098.55  | +2.35    | +58.29   |
| 4.       | 4.  | 4.   | PostgreSQL            | Relational, Multi-model     | 507.19   | +3.82    | +41.08   |
| 5.       | 5.  | 5.   | MongoDB               | Document, Multi-model       | 426.97   | +5.85    | +39.78   |
| 6.       | 6.  | 6.   | IBM Db2               | Relational, Multi-model     | 168.70   | -2.65    | -11.15   |
| 7.       | 7.  |  8. | Elasticsearch         | Search engine, Multi-model  | 151.44   | +1.19    | +8.00    |
| 8.       | 8.  |  7. | Redis                 | Key-value, Multi-model      | 148.75   | +2.51    | -0.27    |
| 9.       | 9.  | 9.   | Microsoft Access   | Relational   | 128.58   | -0.89    | -13.04   |
| 10.      |  11. | 10.  | SQLite                | Relational   | 122.14   | +1.78    | -4.66    |

# IMPEDANCE MISMATCH

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



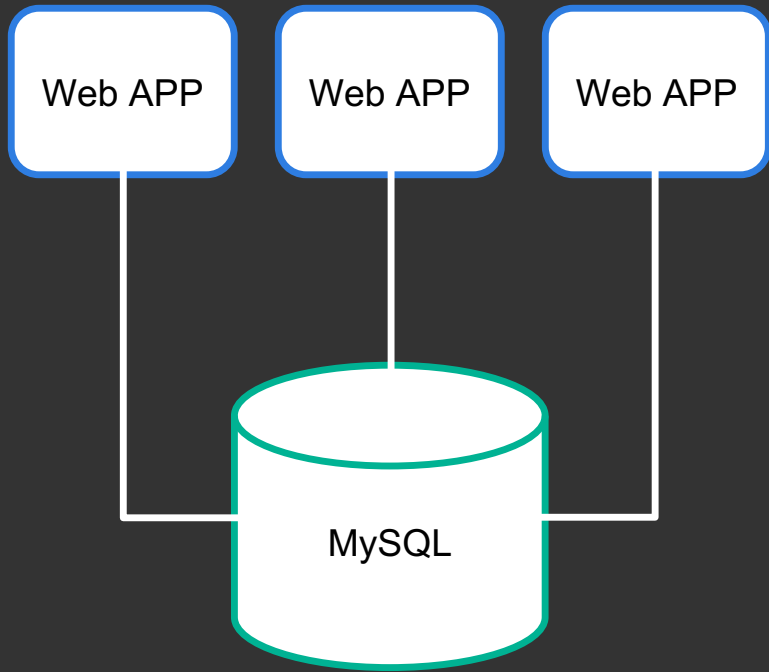






# Travelling Back in Time

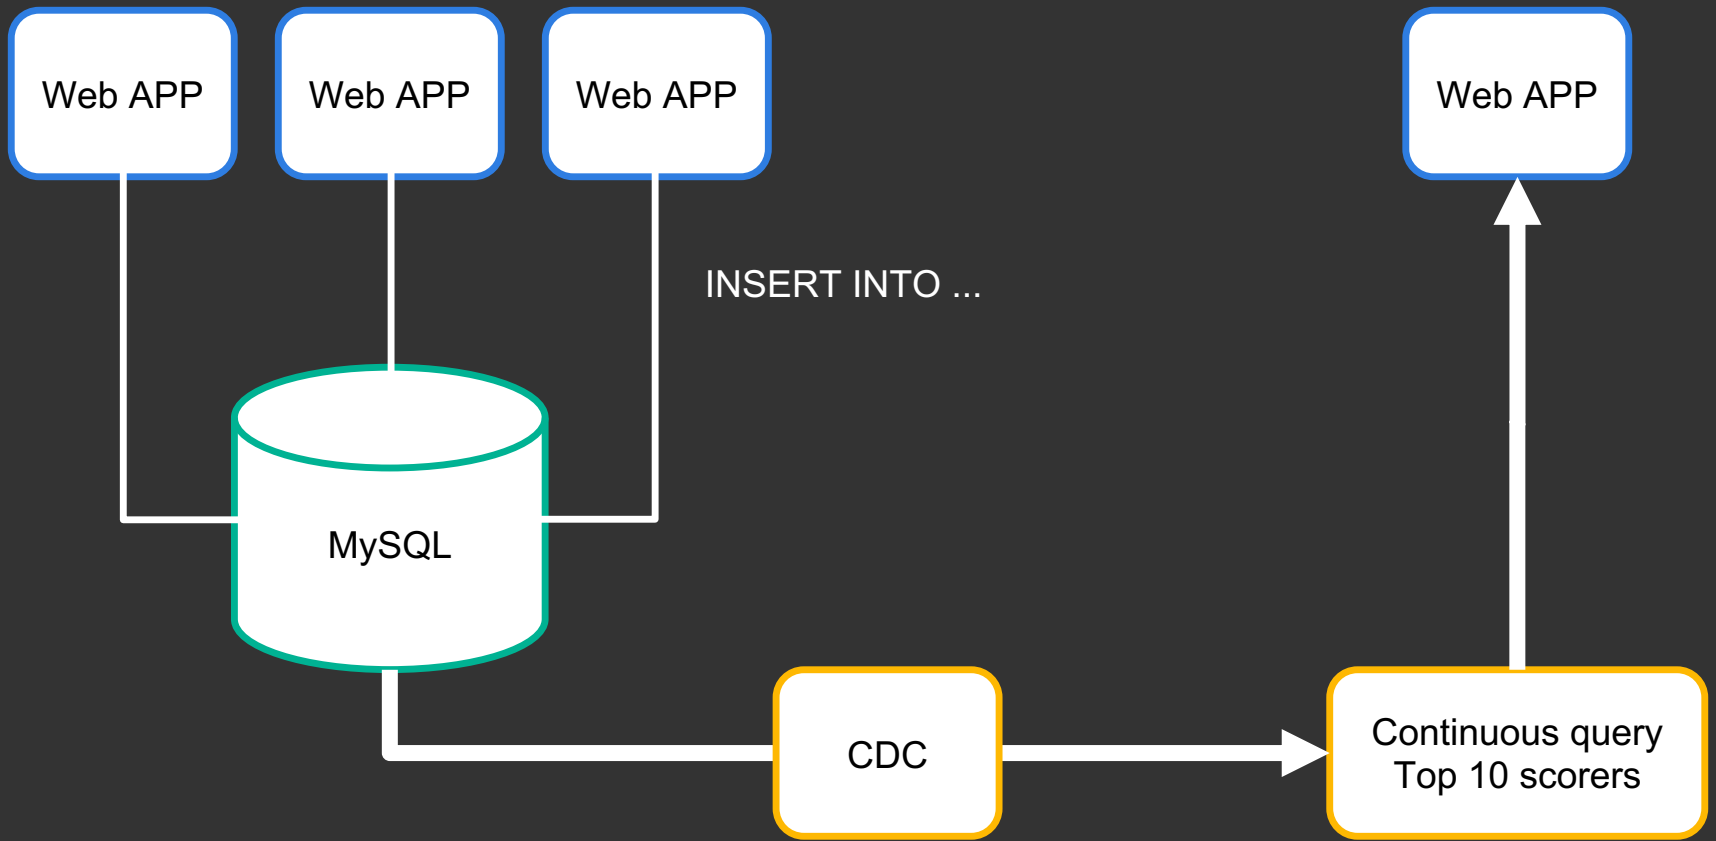
<https://github.com/hazelcast/hazelcast-jet-demos/tree/master/debezium-cdc-without-kafka>



LAMP stack  
GBs of data  
Mostly OLTP,  
OLAP after  
hours



Real-time  
updates =  
Real-time  
expectations



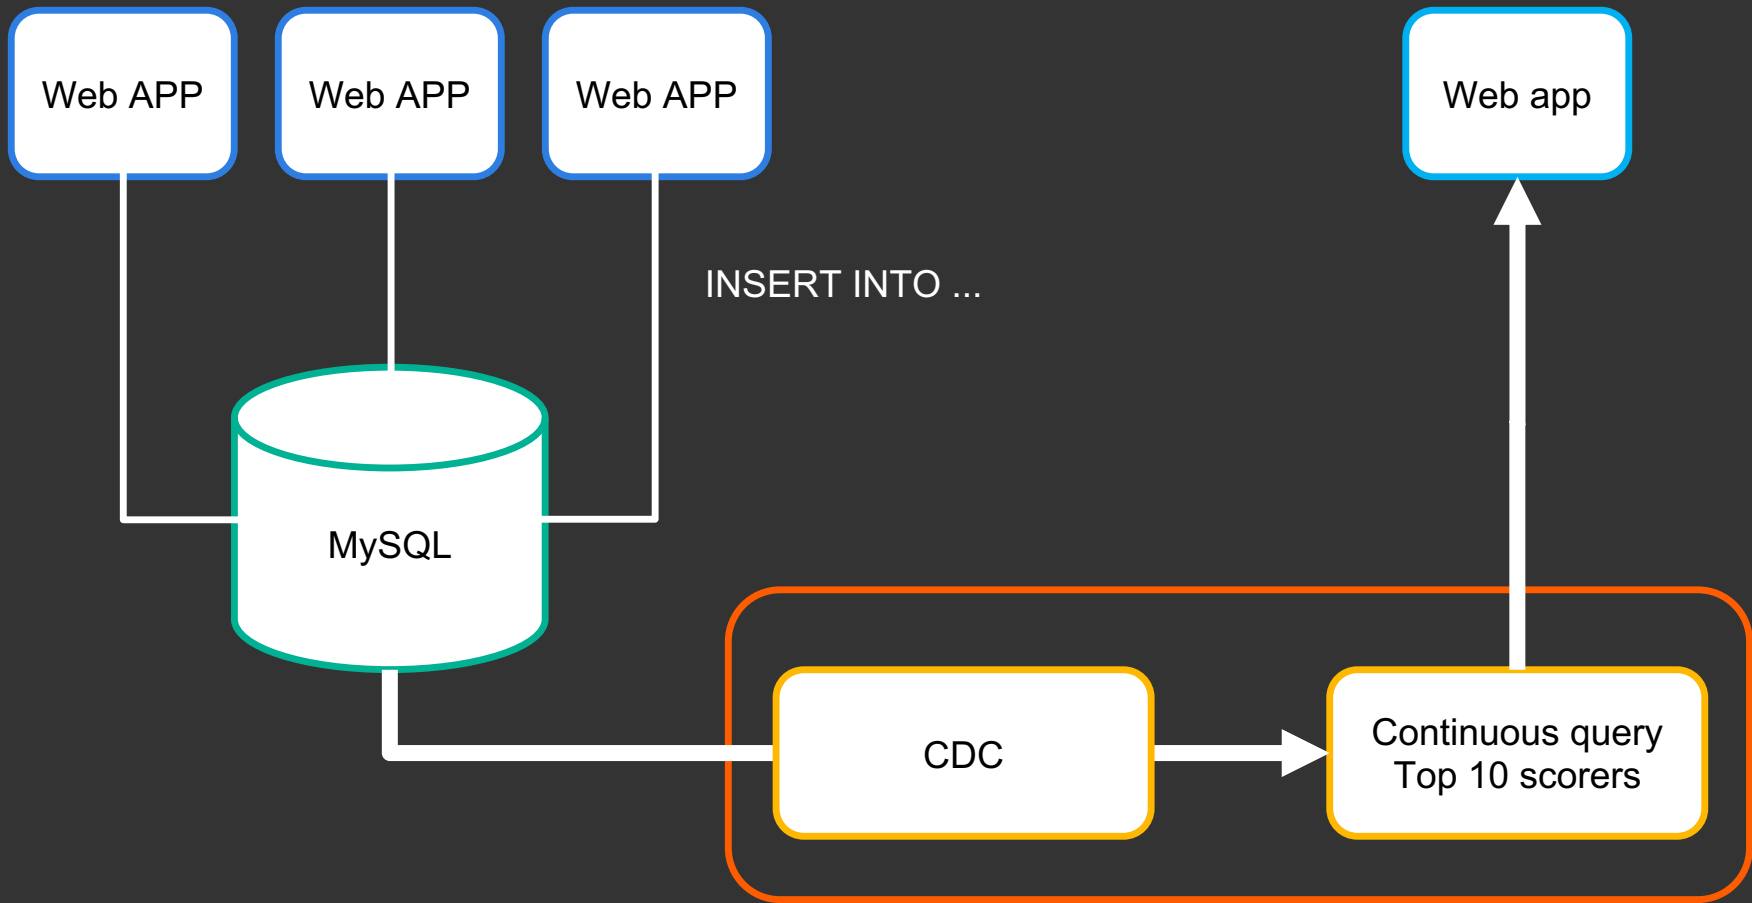
# Hazelcast Jet

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

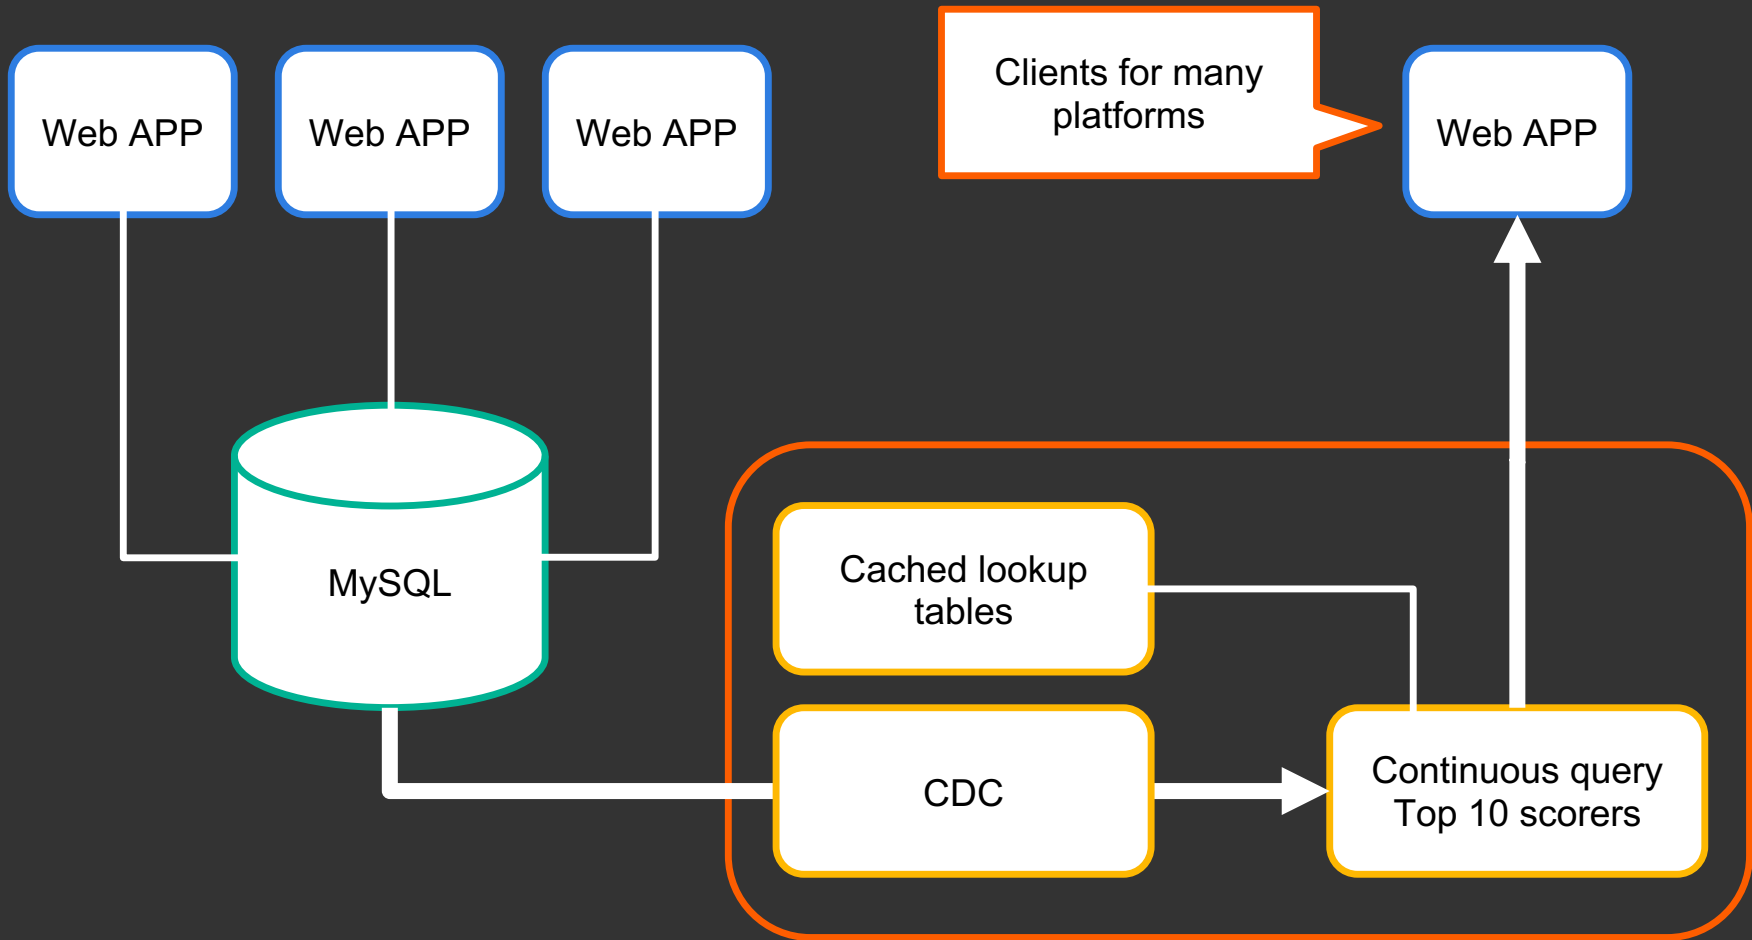
Single Java Binary  
Elastic Clustering  
No ZooKeeper, HDFS..  
Just Java 8 and above

# 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
```

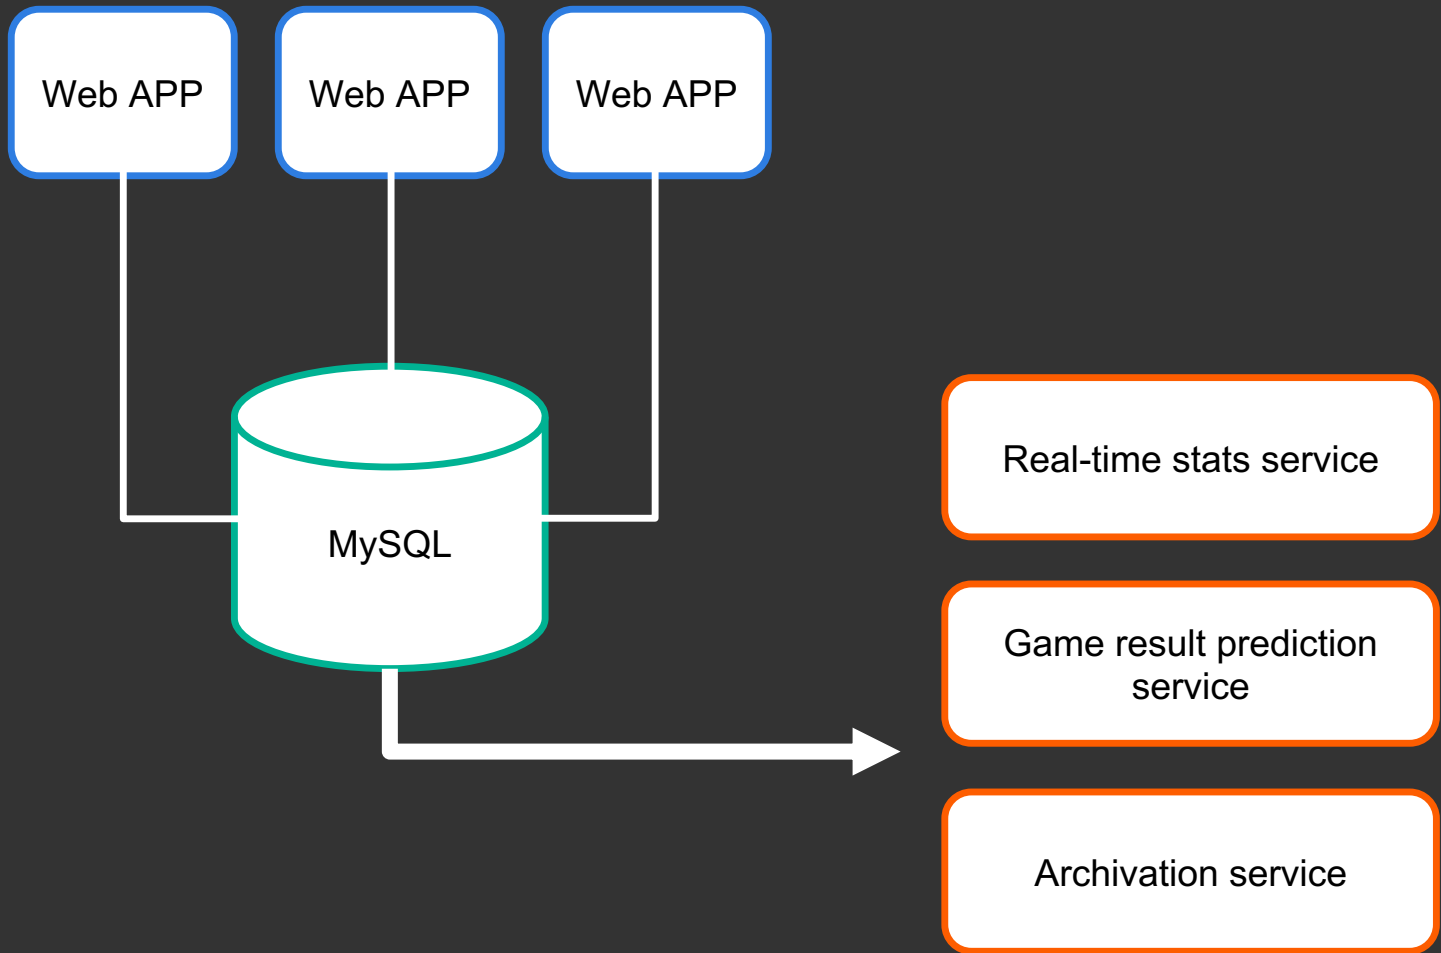


Relational databases usually keep shorter history, compared to dedicated log-based storages.

CDC + SPE + CACHE =

Materialized Views  
Offloaded from the DB

# Modularization Microservices



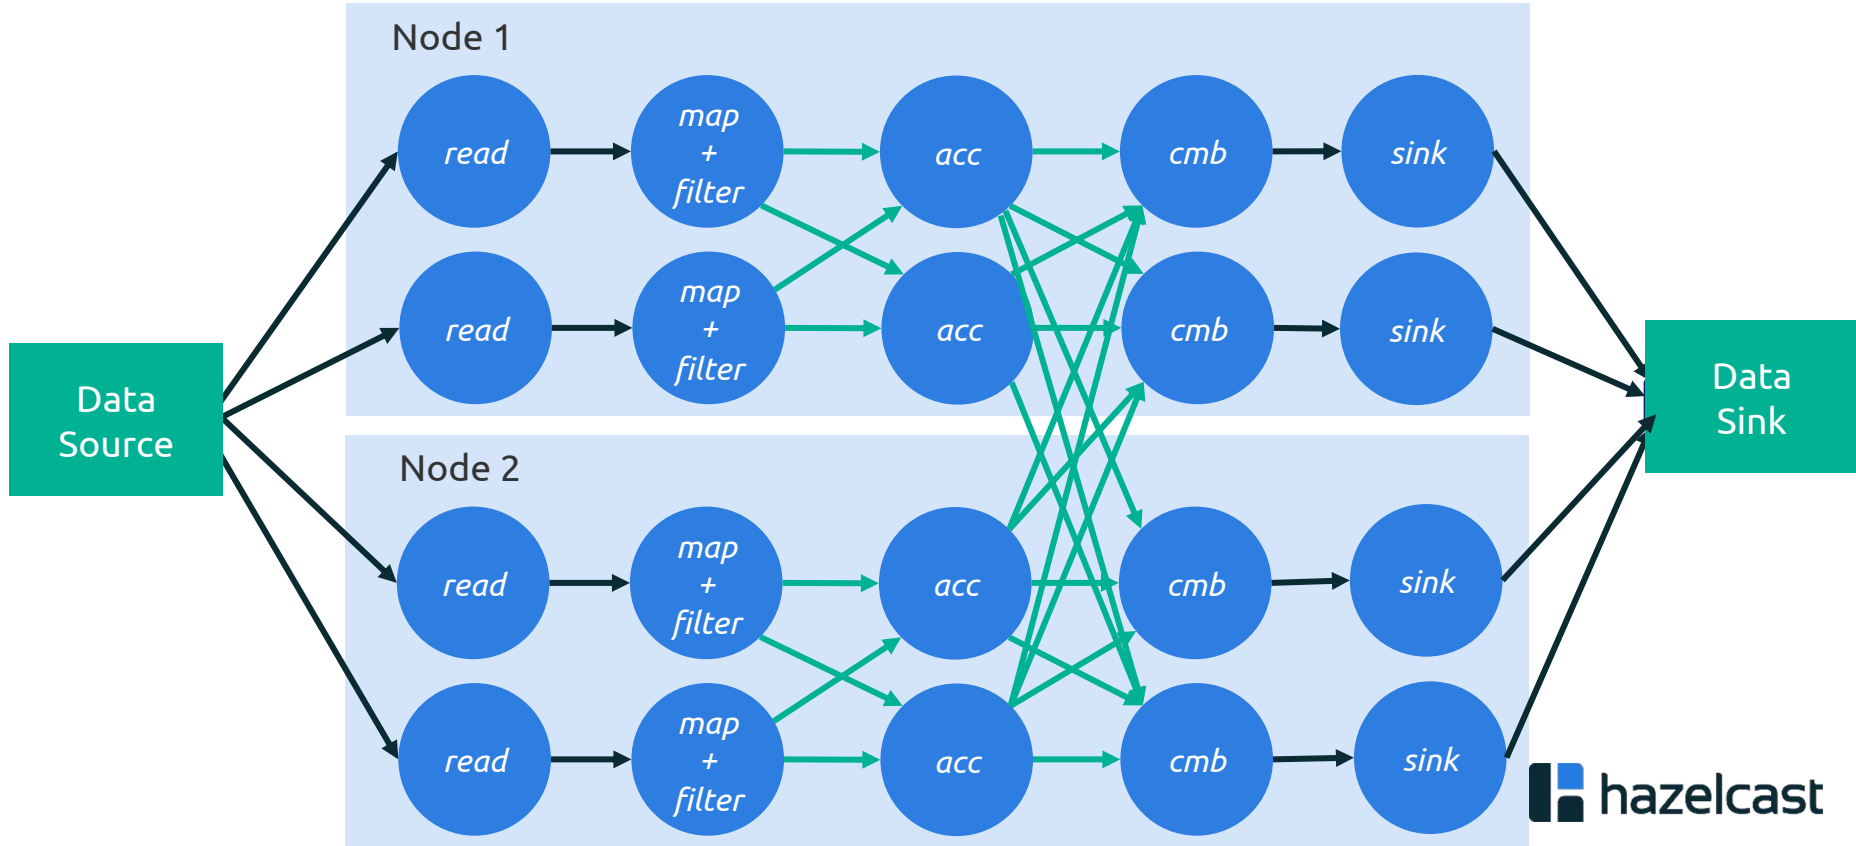
## > Connectors - summary

- Stream processing requires input data presented as streams - sequences of immutable records ("append-only table")
- Streaming frameworks come with clients, agents and connectors (messaging, Kafka, ...)
- CDC - a streaming API to a database.
  - Extracts database changes
  - No standard, many CDC vendors.
- Consider Jet for materialized views in Java



# > Scaling

# > Modern SPEs are build to scale



## > Do I need a streaming cluster?

- A single node can handle 1 million events / s
- Fault-tolerance for instant failover
- Elasticity for performance spikes

## ➤ Fault Tolerance Using Replication

- State of the computation replicated across the cluster
- The tasks of the failed member recovered on other members using the backup replicas
- Regular snapshots for a light-weight F-T
  - Regularly snapshot cluster state and store it reliably
  - Restart computation from last snapshot if it fails
  - Replay a short history of the stream

## › Summary

- Streaming is database "optimised for append-only tables"
- Main UC: Event-driven querying
- Connect streaming to legacy applications using CDC
- Create materialized views outside the database to reduce the database load and modularize your architecture.

<https://jet.hazelcast.org/>

<https://hazelcast.com/blog/how-hazelcast-jet-compares-to-apache-spark/>





# Thank You



519648891000,58094565