

# CockroachDB: From OLTP to HTAP

A Benchmarking Story

*presented by* Arjun Narayan



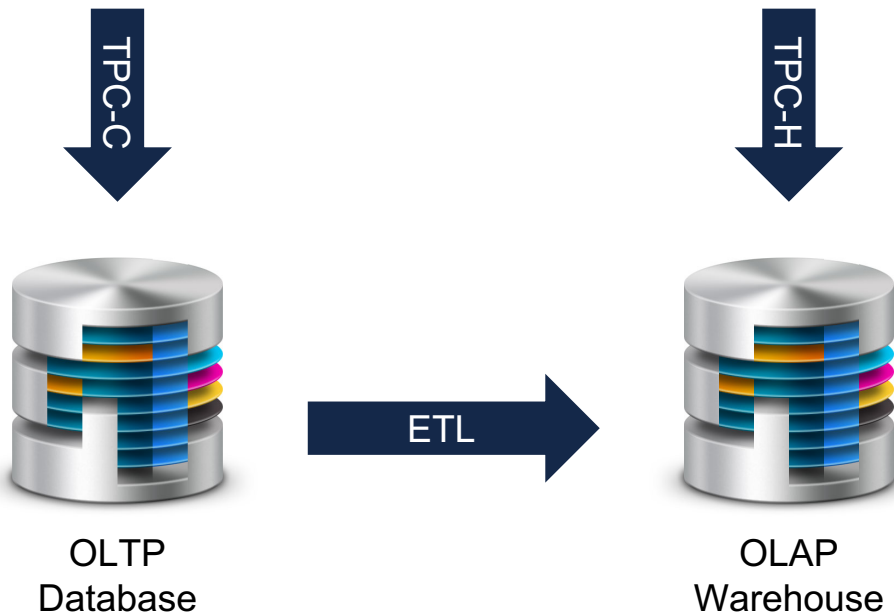
# WARNING!

- Opinionated speculation ahead
- Here's what I want CockroachDB to look like
- Not necessarily what will happen

# HTAP: What is it?

Here's To A Perspective

## Databases today (simplified)




# Databases tomorrow (simplified)

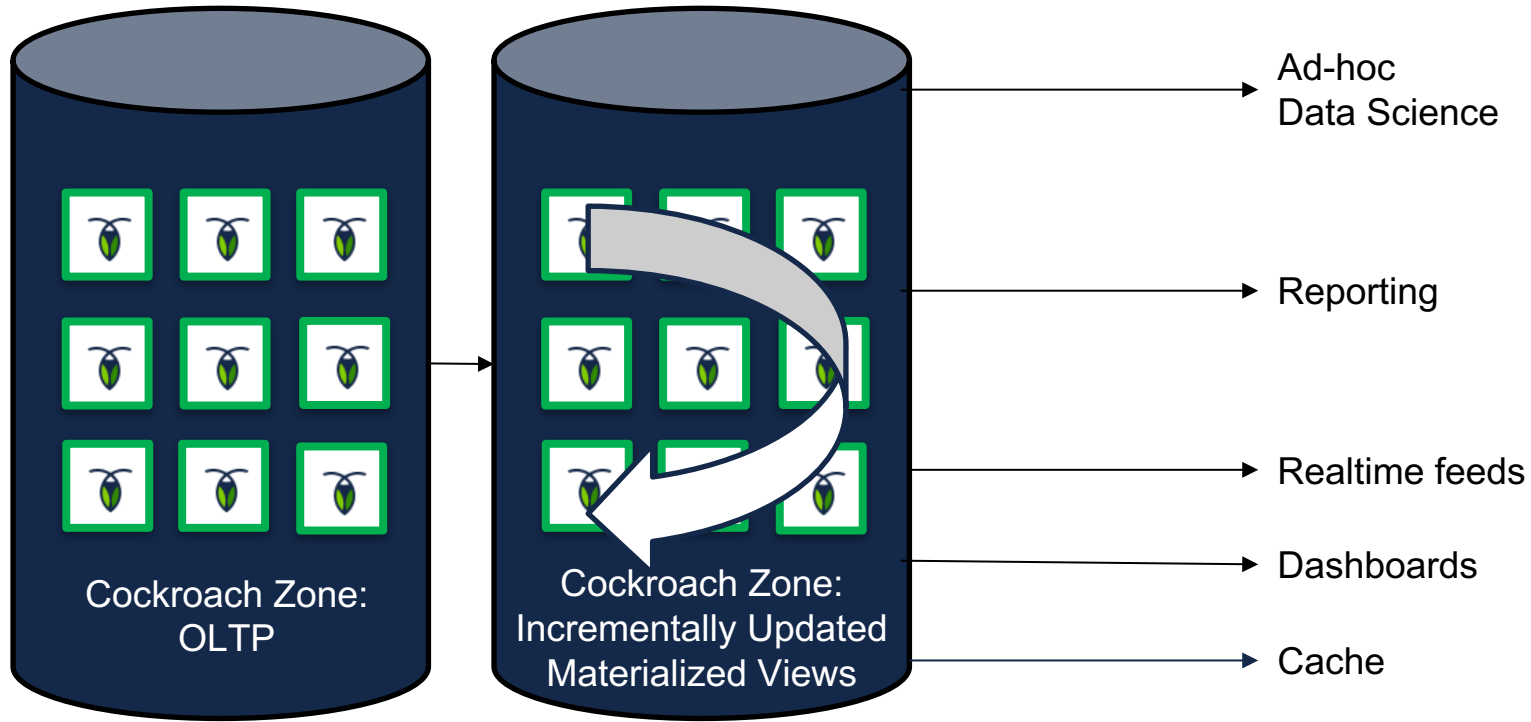


HTAP  
Database

## Data architectures today (reality)



# A vision for data architectures tomorrow



# HTAP: How do we benchmark it?

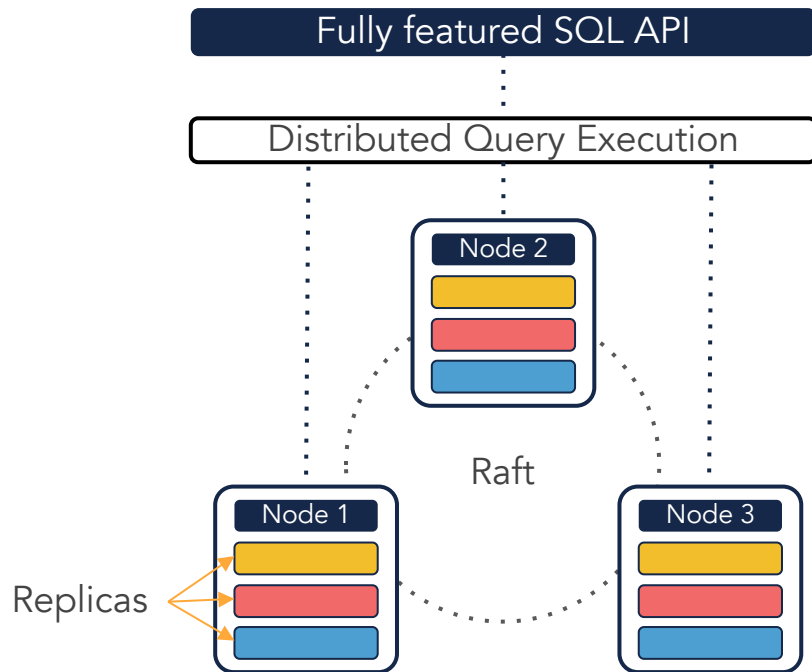
You are what you measure




# CockroachDB Architecture in 60 seconds

Scale out SQL


# Architecture Overview



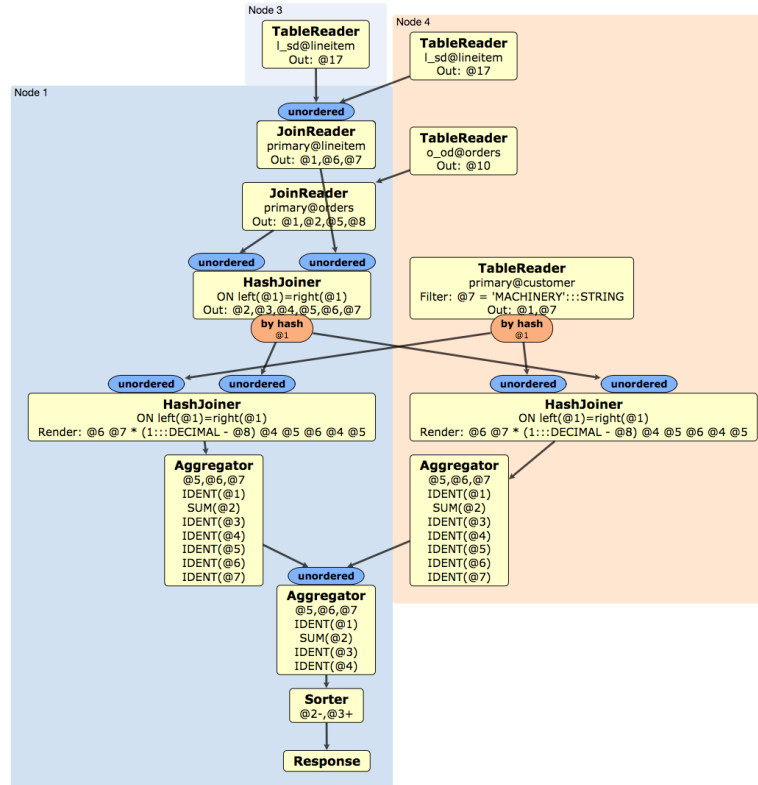
# The CockroachDB KV Layer



# Using Clocks for Fast Reads

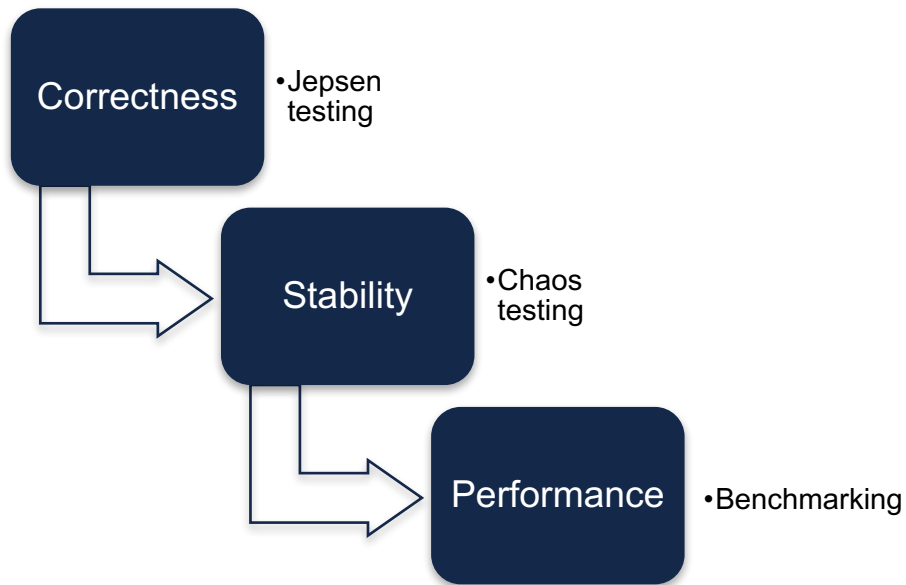


# Distributed Query Execution in CockroachDB

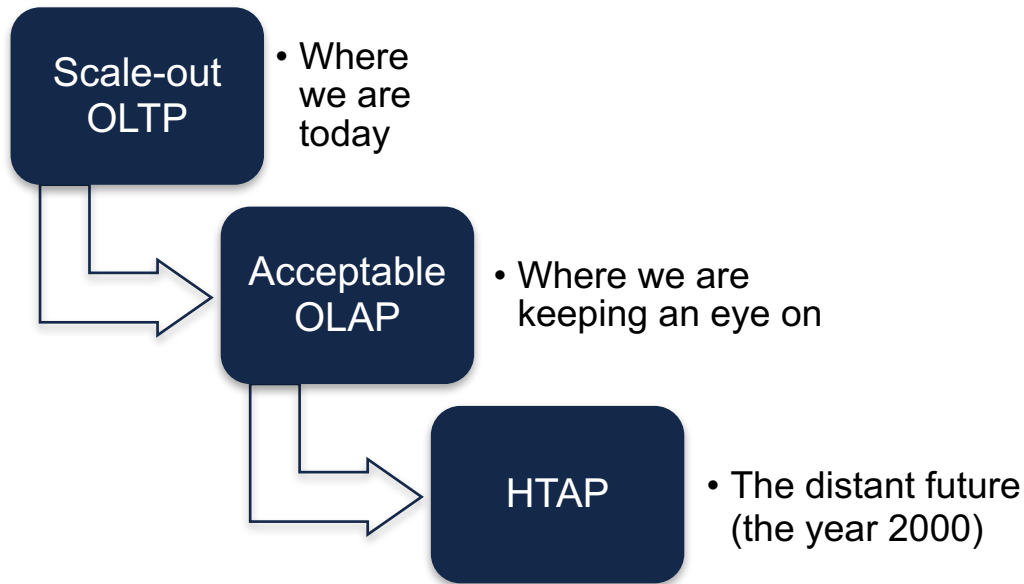


# Building CockroachDB

# Building CockroachDB: An Order of Priorities



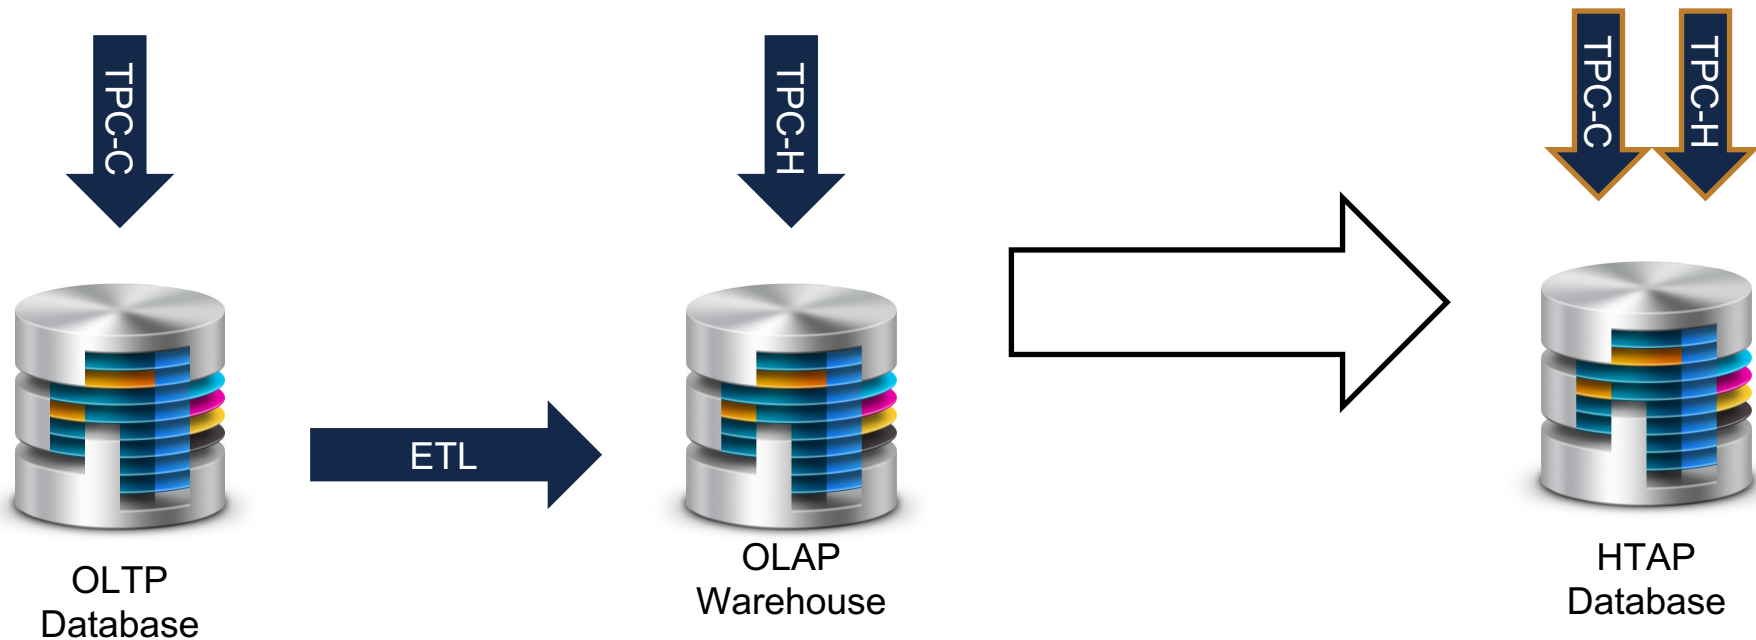
# Building CockroachDB: An Order of Priorities



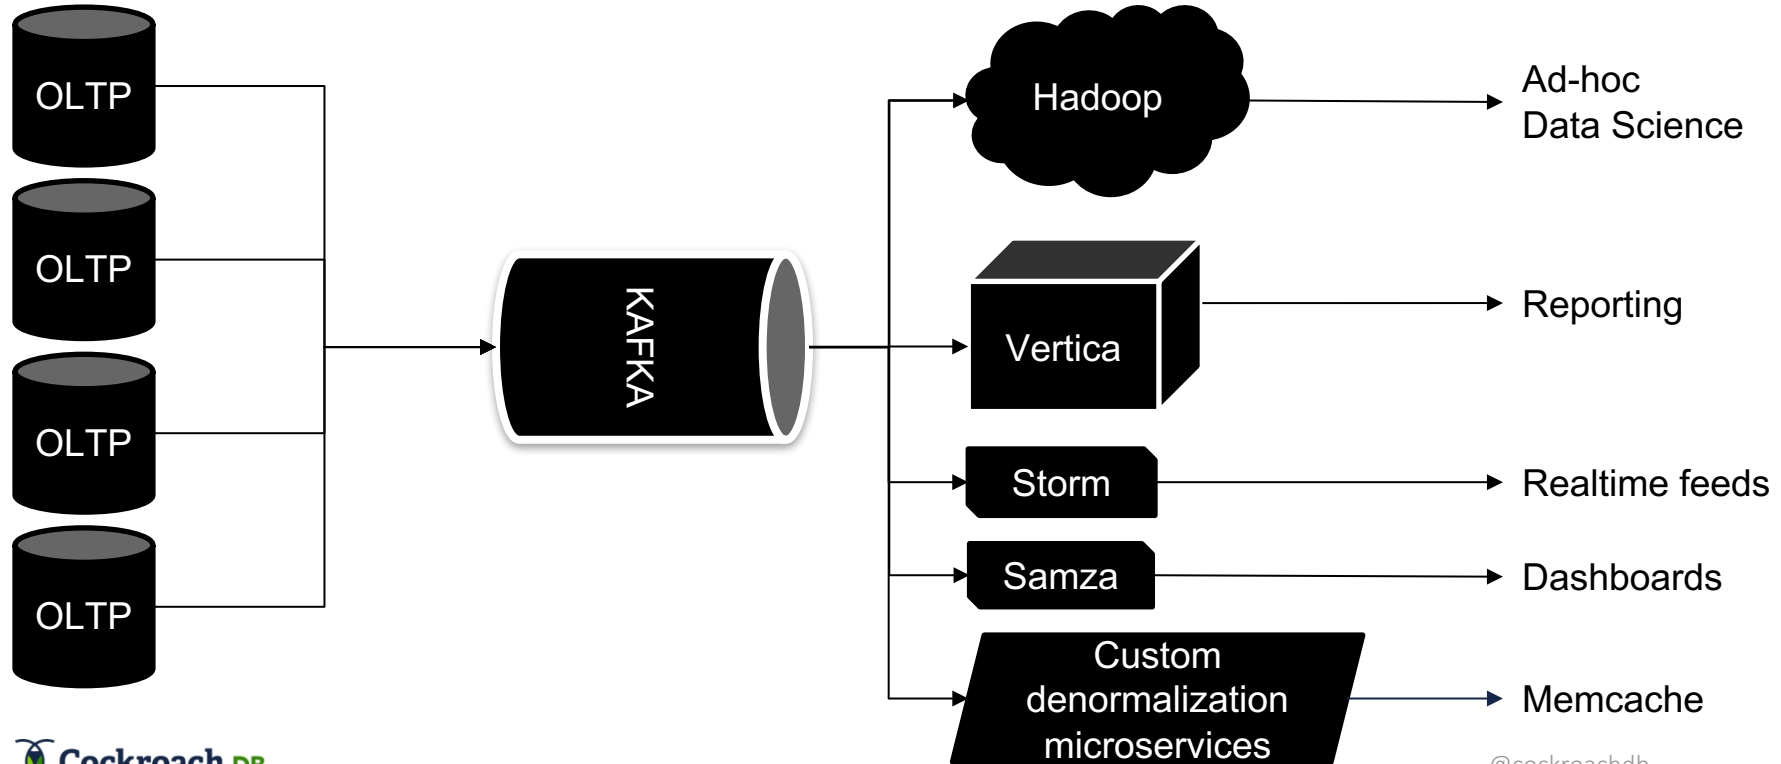


# CockroachDB: An HTAP vision

# HTAP: A strawman



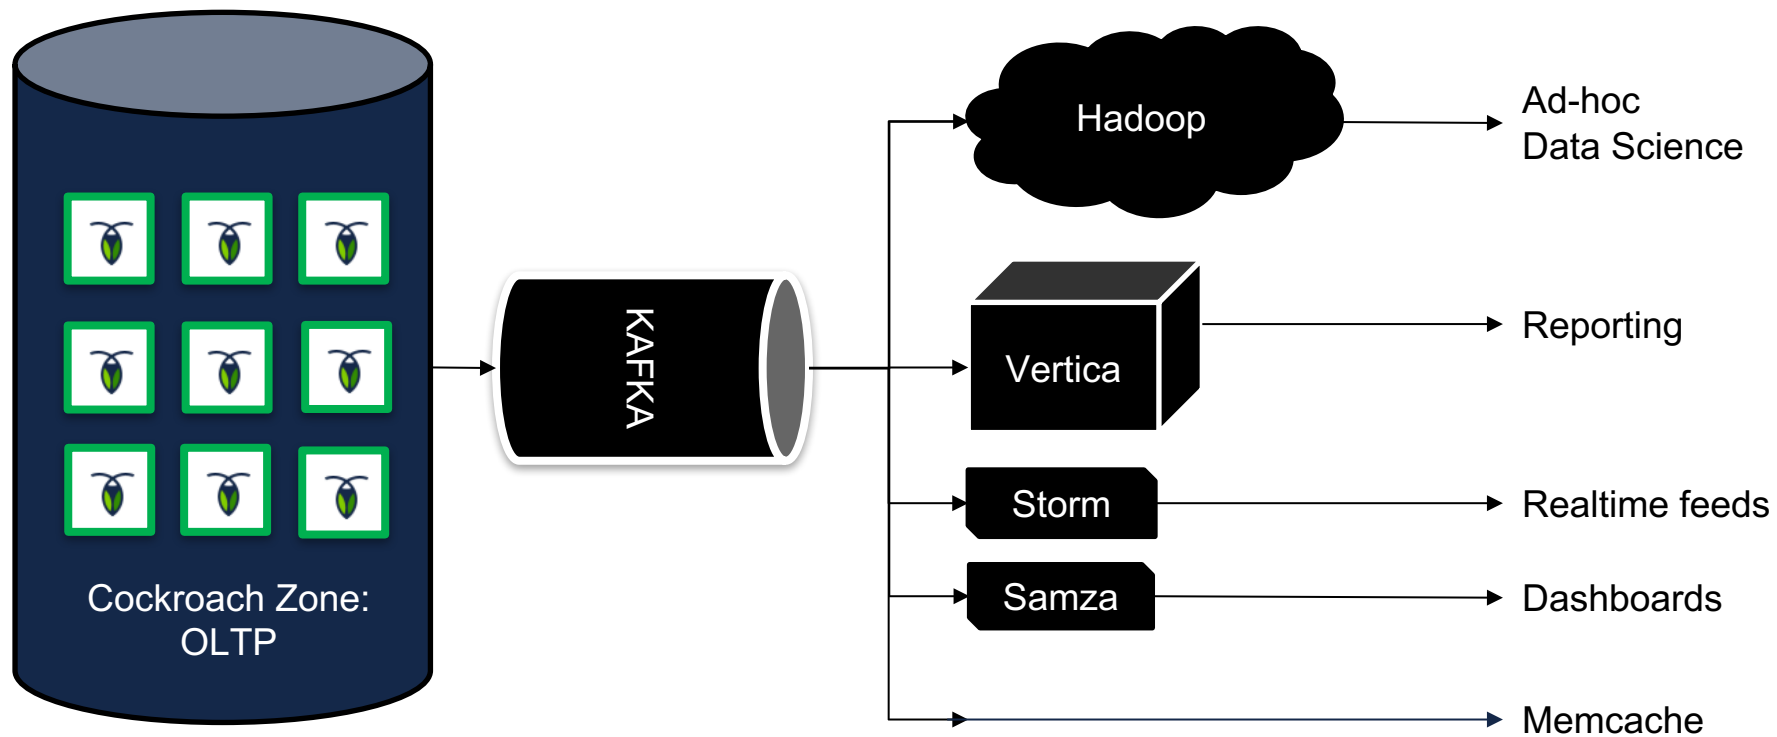
# Data architectures today



# HTAP: Desiderata

- Compute analytics queries and keep results up-to-date
- Resource isolation for OLTP work

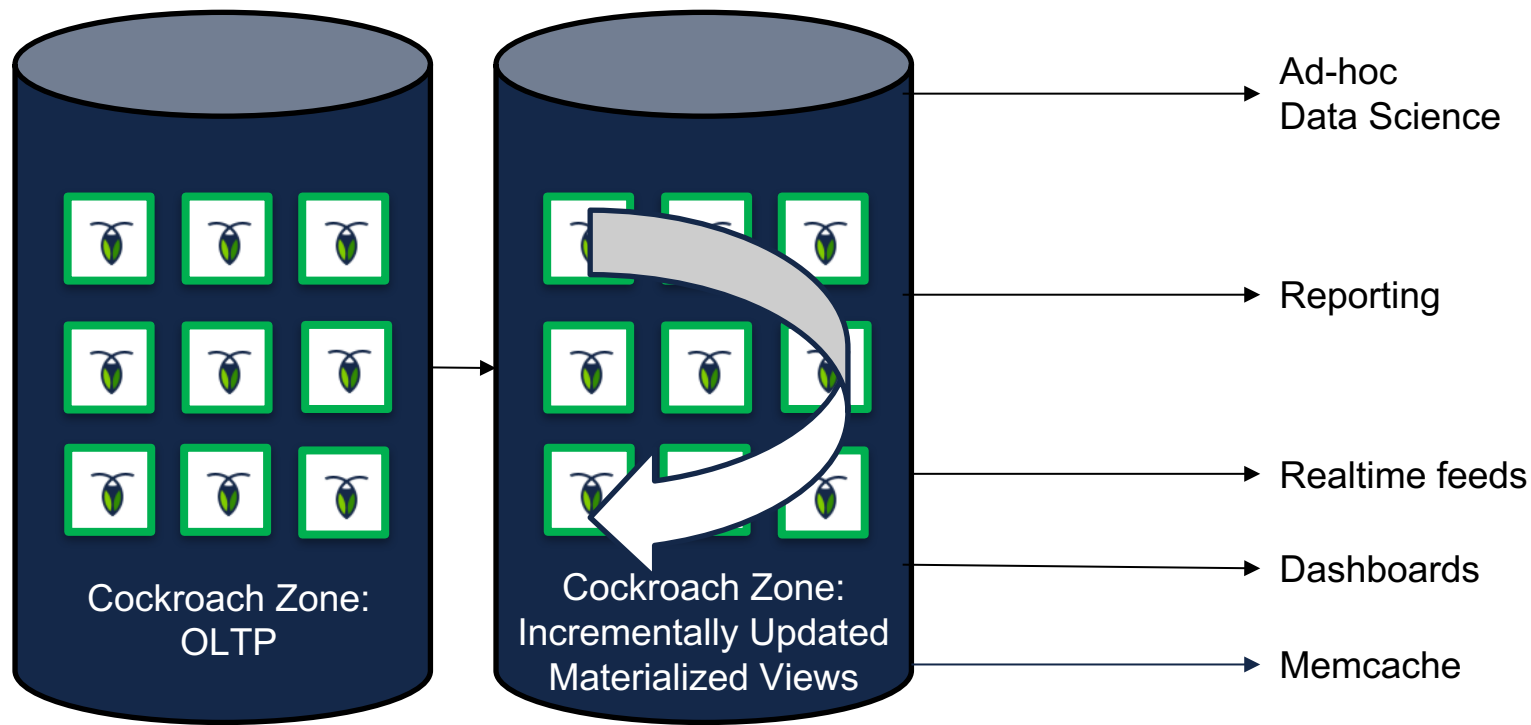
# Data architectures tomorrow



# Incrementally Updated Materialized Views

- Real-time Change feeds from OLTP replicas
- Differential Dataflow for incremental computation of all SQL queries
- Built right into the RDBMS, no separate ETL step
- Cockroach timestamps in dataflow gives serializable transactions that span OLTP+OLAP
- “Read your OLTP writes in OLAP reads” transactional guarantees
- See Naiad (SOSP 2013), Differential Dataflow (CIDR 2013), Frank McSherry’s Blog (GitHub)

## Data architectures the day after



# Benchmarking HTAP: What not to do

- Don't just run TPC-C + TPC-H on the same system
- Why? Materialized views break TPC-H
- It's not just about running ad-hoc TPC-H queries
- HTAP doesn't matter if you can't guarantee predictable TPC-C performance



# Benchmarking materialized views

- How long does it take to compute incremental reports?
- How long does it take to compute ad-hoc queries?
- How long does it take to build a materialized view from scratch?
- How long does it take to propagate OLTP changes?
- How much overhead do we impose on the OLTP system?

# Questions?