

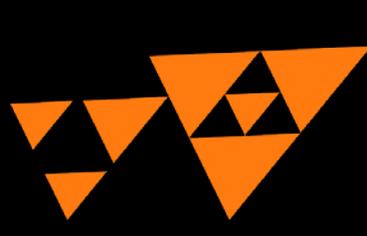
# PELOTON

---

## THE SELF-DRIVING DBMS

@ANDY\_PAVLO / 03 NOV 2016

 H-Store

 **VOLTDB**

2008  
5,000 txn/sec



H-Store: A High-Performance, Distributed Main Memory Transaction Processing System  
*Vldb 2008*

# 2008

5,000 txn/sec

# 2010

11,000 txn/sec



On Predictive Modeling for Optimizing Transaction Execution in Parallel OLTP Systems  
*VLDB 2011*

**2008**

5,000 txn/sec

**2010**

11,000 txn/sec

**2012**

50,000 txn/sec



# 2008

5,000 txn/sec

# 2010

11,000 txn/sec

# 2012

50,000 txn/sec

# 2015

4,000,000 txn/sec



TicToc: Time Traveling Optimistic Concurrency Control  
*SIGMOD 2016*



ONLY MAXIMIZING OLTP  
THROUGHPUT LEADS TO AN  
**UNFULFILLING** LIFE.



Got a tip? [Let us know.](#)

Follow Us [f](#) [@](#) [v](#) [R](#) [in](#) [g+](#) [RSS](#)

[News](#) [Video](#) [Events](#) [Crunchbase](#)

Search

**HOT NEWS** Wesley Chow voted 'sexiest' CTO of New York ▶

Apple

Omniture

foundationdb

Fundings & Exits

Popular Posts

# Apple Acquires Durable Database Company FoundationDB

Posted Mar 24, 2015 by [Matthew Panzarino \(@panzer\)](#)

4,790 SHARES [f](#) [t](#) [in](#) [g+](#) [v](#) [R](#) [in](#) [g+](#) [RSS](#) [Email](#) [Print](#)



Apple has acquired [FoundationDB](#), a company that specializes in speedy, durable NoSQL databases, TechCrunch has learned.

## CrunchBase

**Apple** -

FOUNDED  
1976

### OVERVIEW

Apple is a multinational corporation that designs, manufactures, and markets mobile communication and media devices, personal computers, portable digital music players, and sells a variety of related software, services, peripherals, networking solutions, and third-party digital content and applications. Apple provides many products and services, including iPhone; iPad; iPod; Mac; iPod; Apple TV; a ...

### LOCATION

[Cupertino, California](#)

### CATEGORIES

[Computers, Electronics](#), [Consumer Electronics](#), [Hardware + Software](#), [Retail](#)

### WEBSITE

<http://www.apple.com>

[Full profile for Apple](#)

▲ Gurrewe 19 days ago

The company I'm working for is looking for a way to scale our DB-layer. Anyway, FoundationDB was more or less the only candidate against MySQL.

Does anyone know of any good (and proven) alternatives to FoundationDB?

▲ tracker1 19 days ago

It's funny you mention that.. but actually hiring a part-time PostgreSQL DBA is all but impossible, I reached out to most of the support companies listed on the north american website... mainly I wanted for someone to setup a small (3-node) replica set of the most recent version of postgres with plv8 some sane backup scripts and pretty much nobody replied... EnterpriseDB won't talk to you without laying out at least \$10k to start, and I would rather pay a person (or small company) I can call that to get things running... more if it kept running well.

AVERAGE SALARY FOR  
DATABASE ADMINS IN  
2015 WAS \$81,710.

# SELF-DRIVING

A DBMS THAT CAN CONFIGURE,  
TUNE, AND OPTIMIZE ITSELF  
WITHOUT ANY HUMAN  
INTERVENTION.

# YES

DATABASE DESIGN

DATA PLACEMENT

QUERY OPTIMIZATION

KNOB CONFIGURATION

BACK-UP & RECOVERY

PROVISIONING

# NO

SECURITY & ACLs

DATA INTEGRATION

UNPLANNED HALTS

VERSION CONTROL

# WHAT'S NEW?

PREVIOUS EFFORTS ARE  
REACTIVE & HUMAN-DRIVEN.

A SELF-DRIVING DBMS HAS TO  
BE PREDICTIVE.

# WHY NOW?

RECENT ADVANCEMENTS IN  
HARDWARE AND DEEP NEURAL  
NETWORKS MAKE AUTONOMOUS  
OPERATION NOW POSSIBLE.



Peloton



Non-volatile ↘

IN-MEMORY

OLTP+OLAP

LLVM EXEC

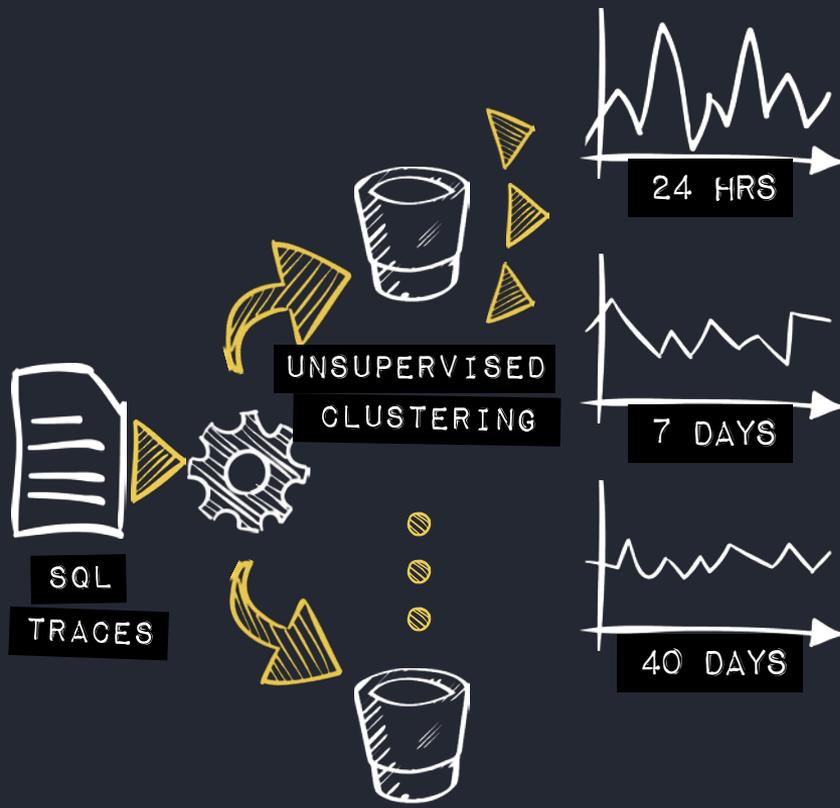
AUTONOMOUS

# THE BRAIN

INTEGRATED DEEP LEARNING  
FRAMEWORK TO MODEL,  
PREDICT, AND OPTIMIZE HTAP  
DATABASE WORKLOADS.

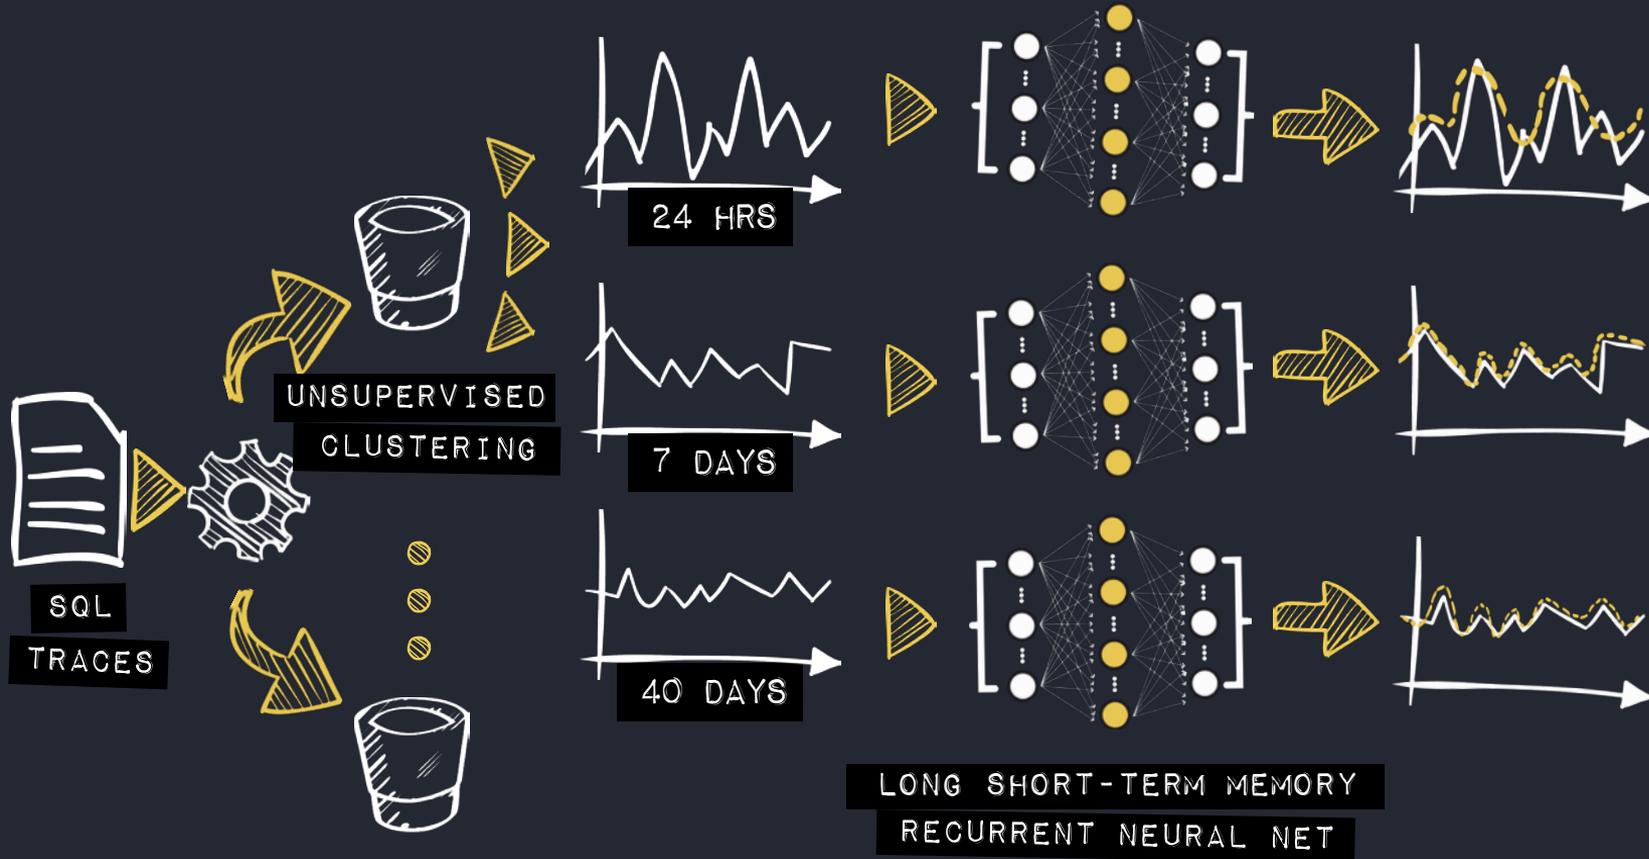


# WORKLOAD CATEGORIZATION



# WORKLOAD CATEGORIZATION

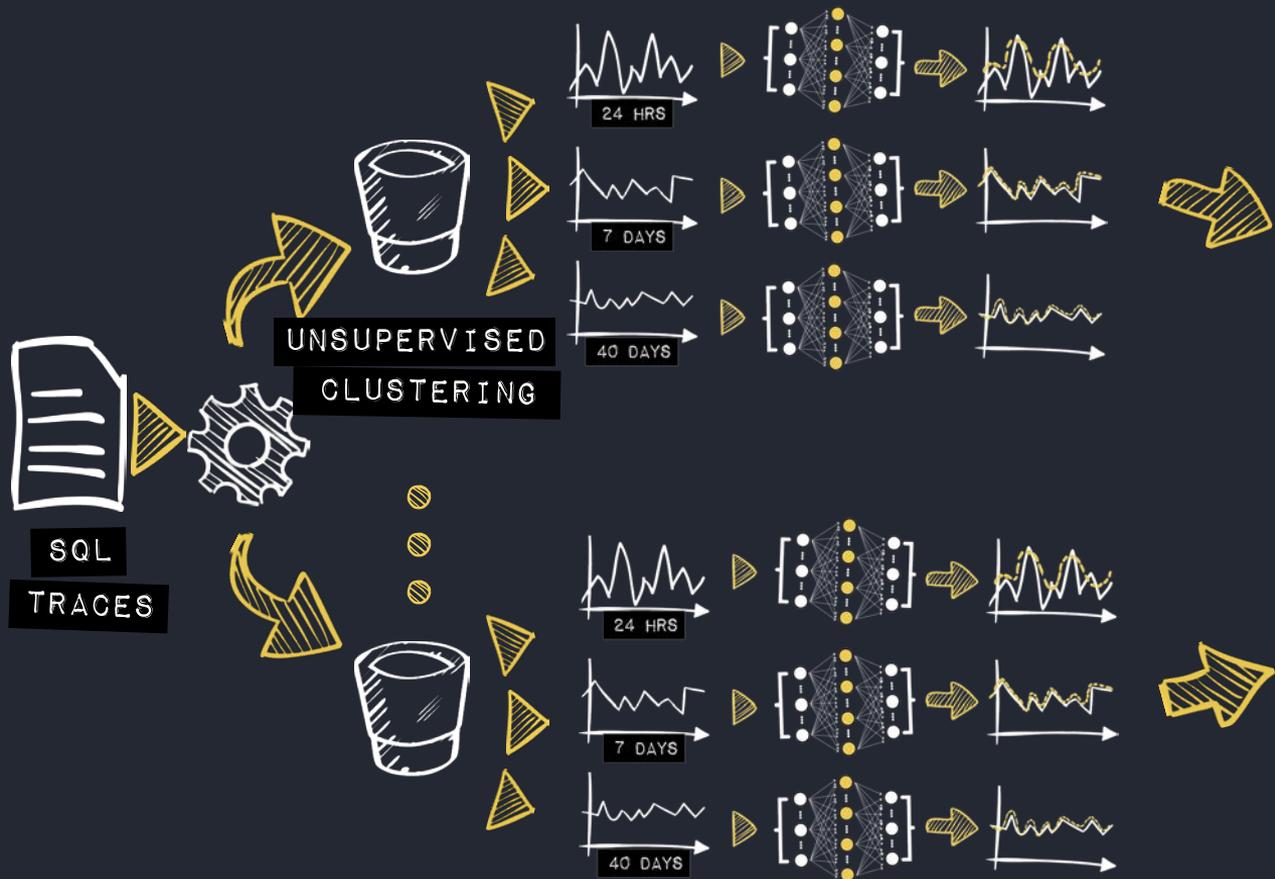
# WORKLOAD FORECASTING



# WORKLOAD CATEGORIZATION

# WORKLOAD FORECASTING

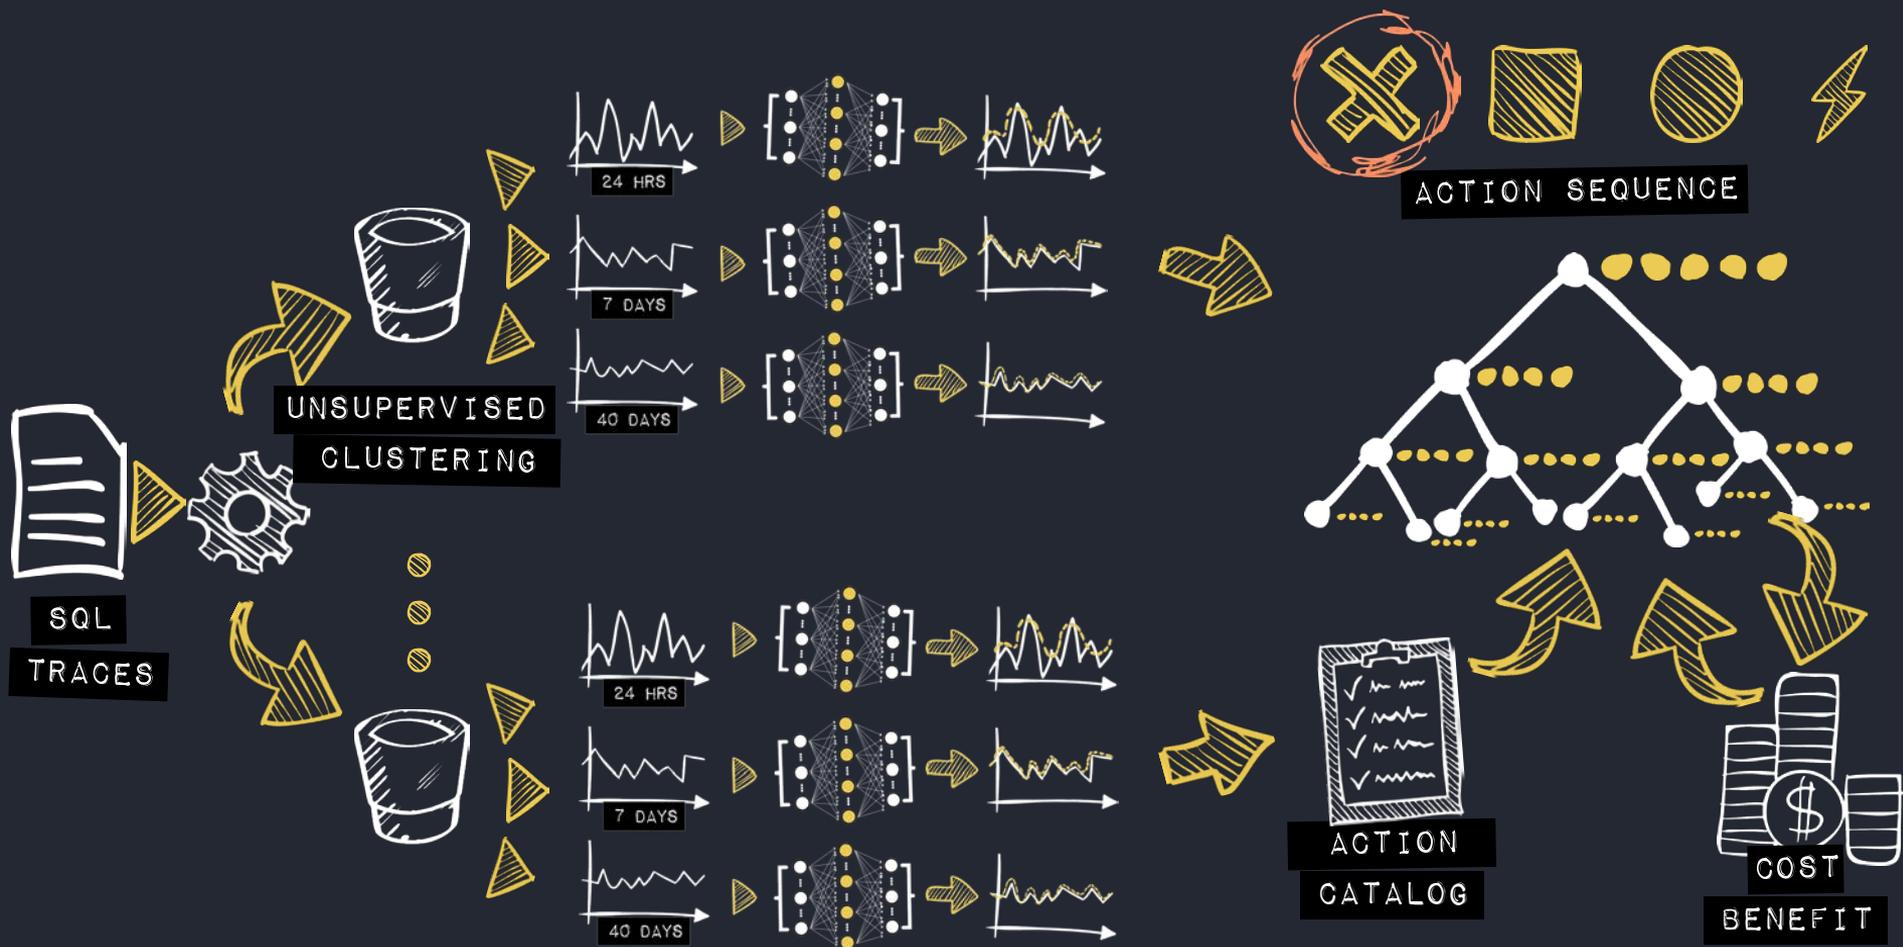
# OPTIMIZATION PLANNING



# WORKLOAD CATEGORIZATION

# WORKLOAD FORECASTING

# OPTIMIZATION PLANNING



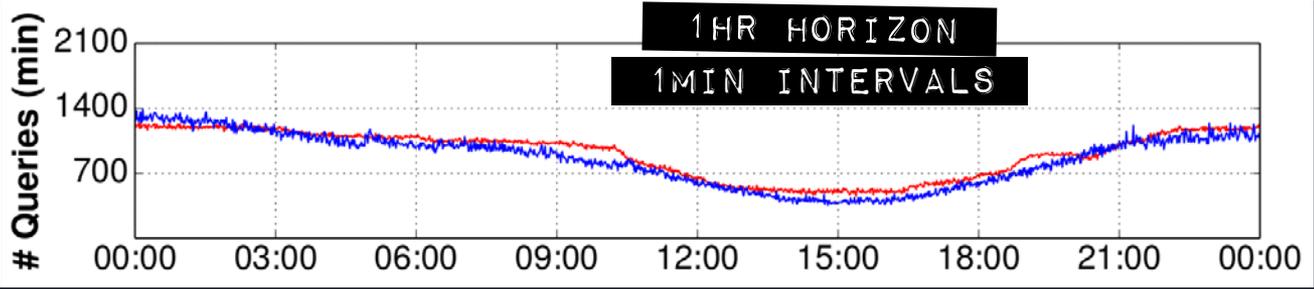
# EVALUATION

SYNTHETIC WORKLOAD BASED  
ON REDDIT TRAFFIC DATA.

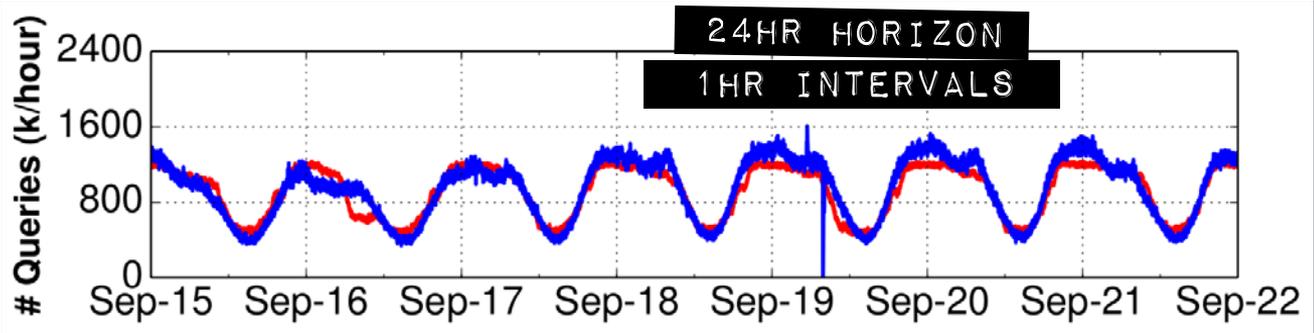
FORECAST WITH TENSORFLOW.

ADAPTIVE STORAGE.

— Actual Workload    — Predicted Workload



**ERROR RATE: 14.7%**  
**CPU TRAINING: 25MIN**  
**PROBE: 2MS**  
**UPDATE: 5MS**  
**SIZE: 2MB**



**ERROR RATE: 17.9%**  
**CPU TRAINING: 18MIN**  
**PROBE: 2MS**  
**UPDATE: 5MS**  
**SIZE: 2MB**

# ADAPTIVE STORAGE

CHANGE THE LAYOUT OF DATA  
OVER TIME BASED ON HOW IT  
IS ACCESSED.



Bridging the Archipelago Between Row-Stores and  
Column-Stores for Hybrid Workloads

*SIGMOD 2016*

OLTP



```
UPDATE myTable
  SET A = 123,
      B = 456,
      C = 789
  WHERE D = "xxx"
```

OLAP



```
SELECT AVG(B)
  FROM myTable
  WHERE C < "yyy"
```

ORIGINAL TABLE

A	B	C	D

Hot

Cold

OLTP



```
UPDATE myTable
SET A = 123,
    B = 456,
    C = 789
WHERE D = "xxx"
```

OLAP



```
SELECT AVG(B)
FROM myTable
WHERE C < "yyy"
```

ORIGINAL TABLE

Hot

A	B	C	D

Cold

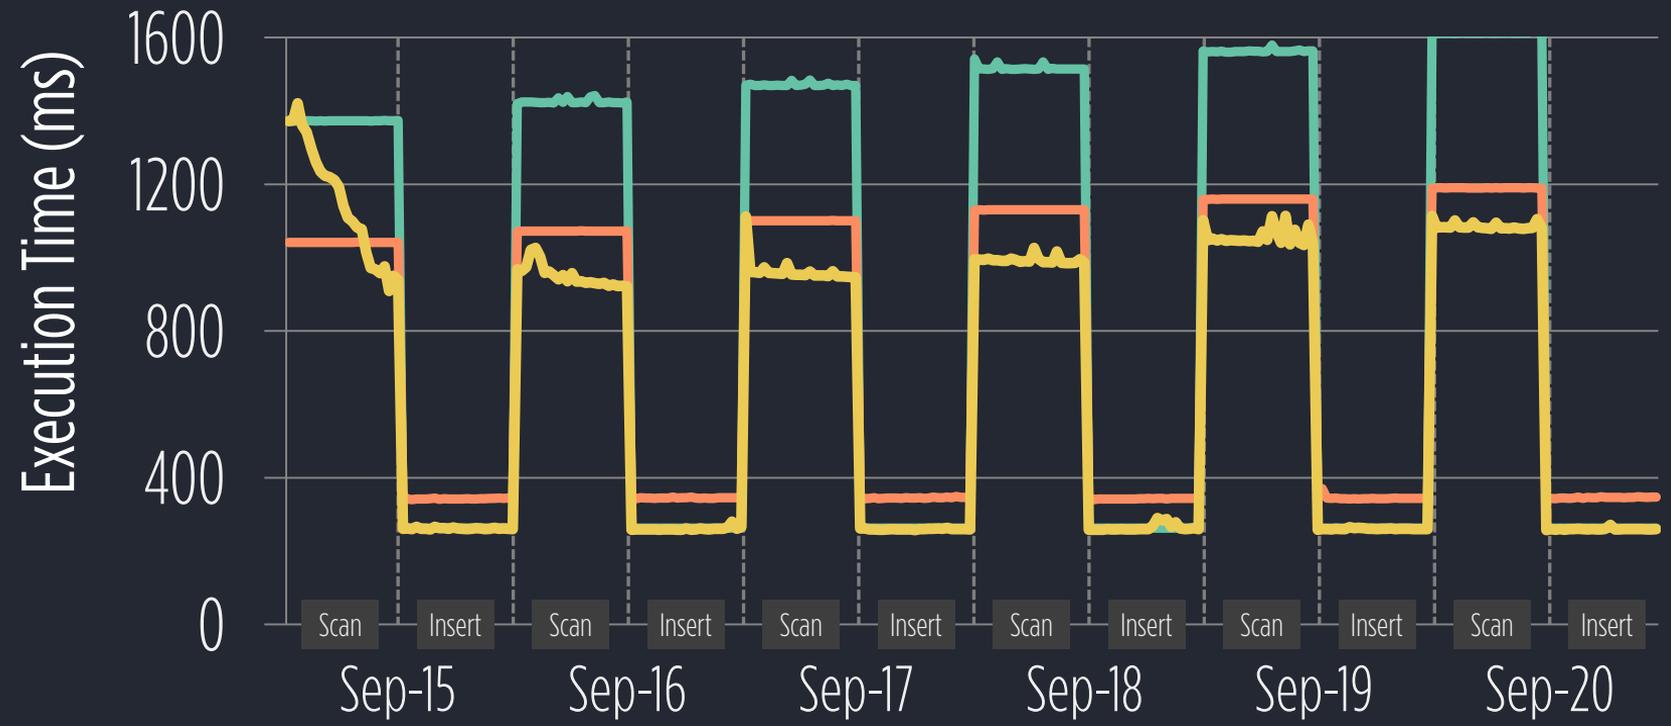
ADAPTED TABLE

A	B	C	D

A	B	C	D



— Row Layout — Column Layout — Adaptive Layout



# CURRENT STATUS

SINGLE-NODE ONLY.

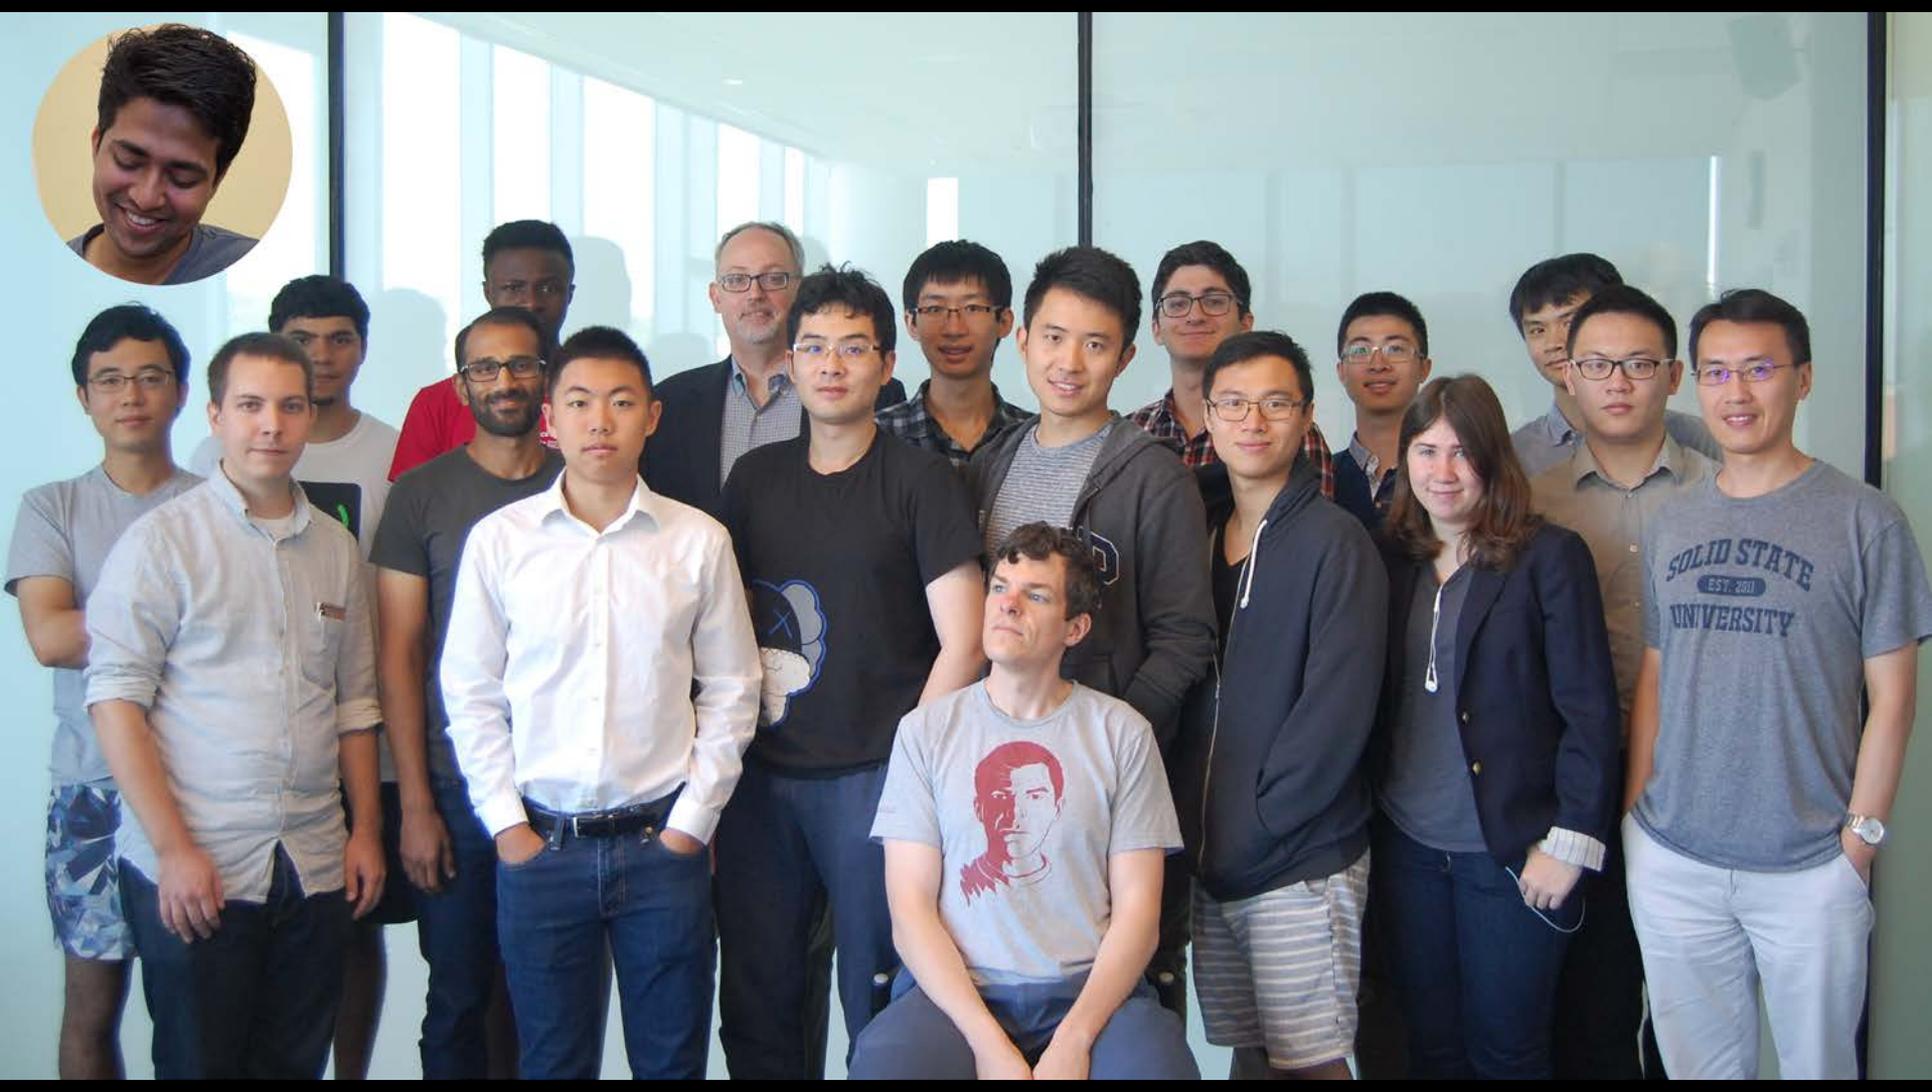
LOOKING FOR REAL-WORLD  
DEPLOYMENTS TO TEST.

APACHE 2.0 LICENSE



# Peloton

<http://pelotondb.org>



END

@ANDY\_PAVLO