CARNEGIE MELLON
DATABASE GROUP

I DON'T WANT TO BE THE
# MITT ROMNEY
OF DATABASES

@andy_pavlo

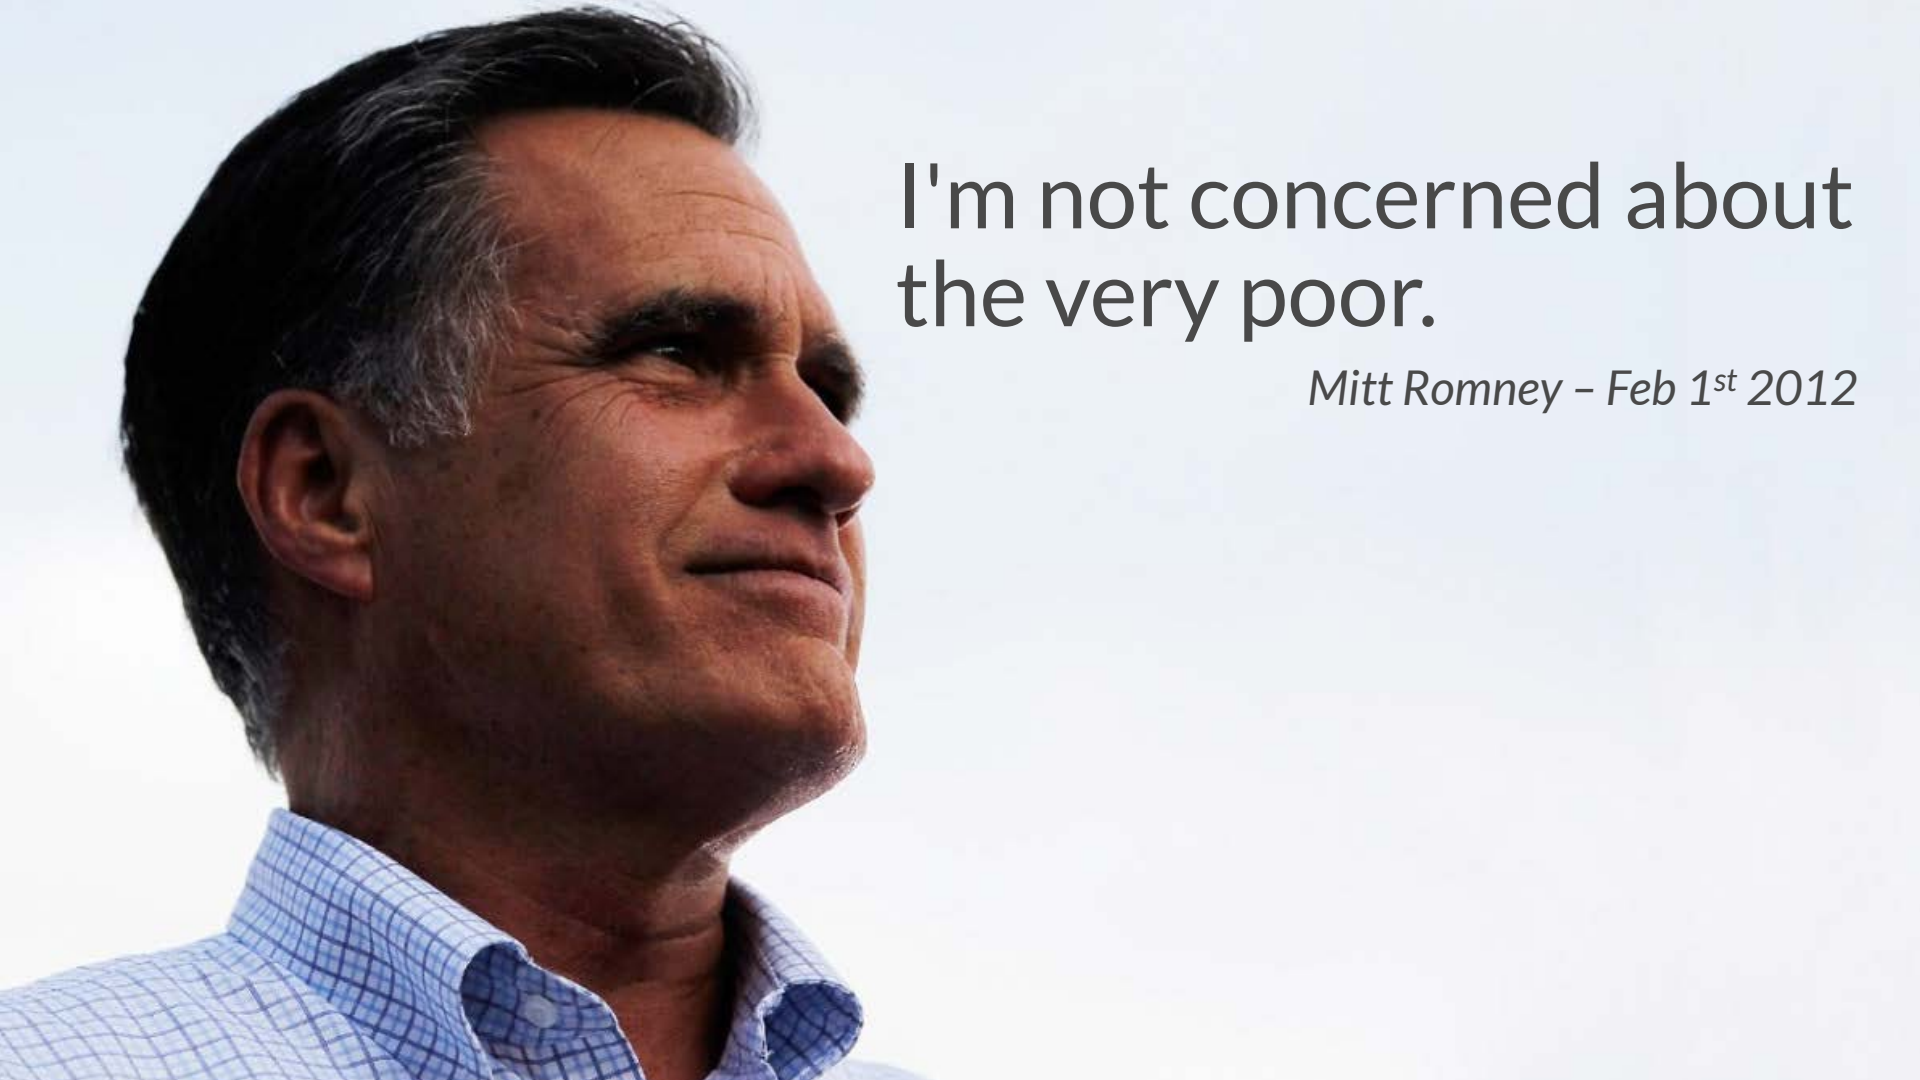
I'm not concerned about the very poor.
*Mitt Romney – Feb 1st 2012*

# 1%

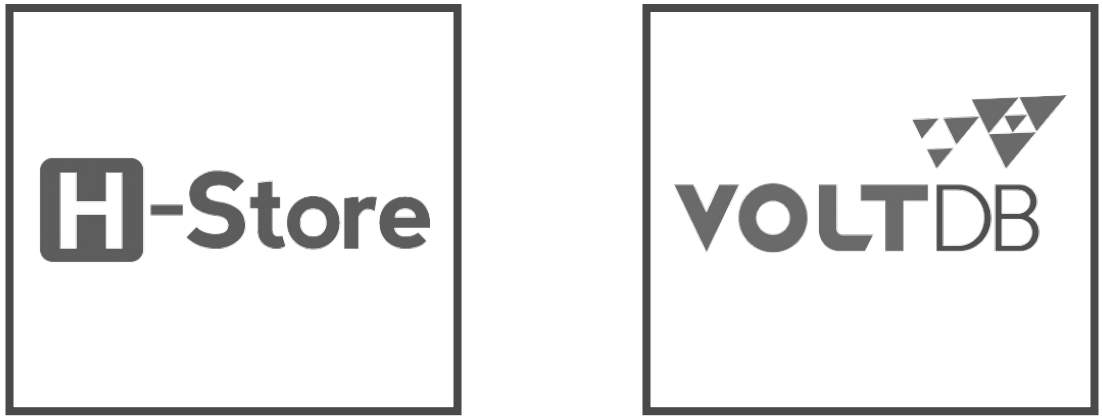


**(a)** H-Store **(b)** VoltDB

**Figure 1:** Examples of database systems for the “one percent.”

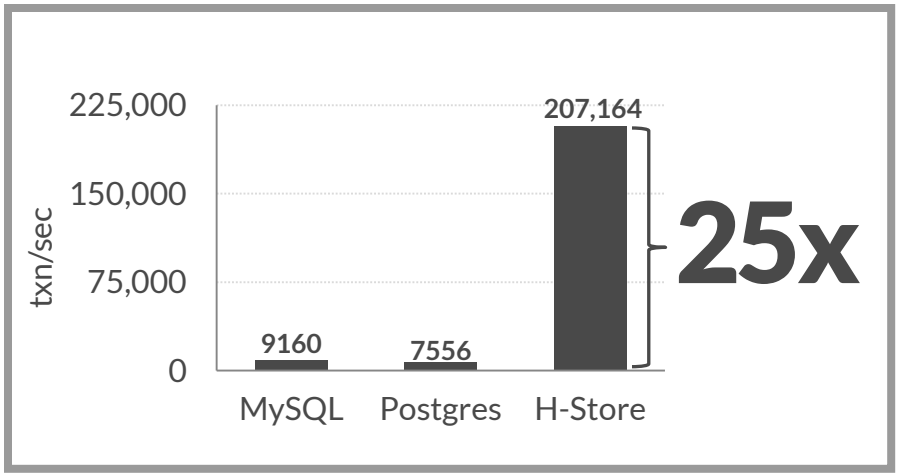


**Figure 2:** Throughput of three DBMSs for the single-node Voter benchmark with 100% single-partition transactions.

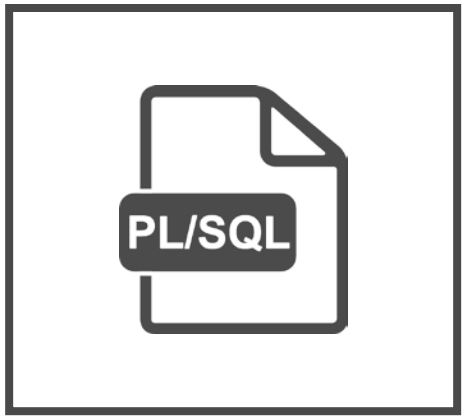


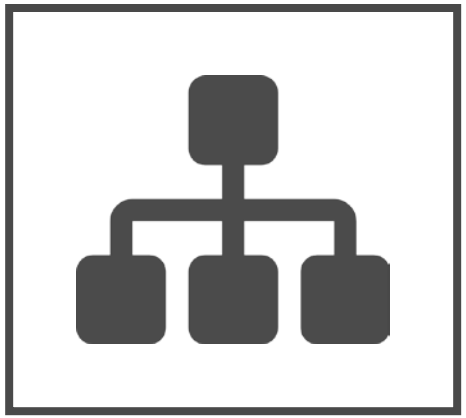

**(a)** Stored Procedures **(b)** Separate OLAP DBMS **(c)** Pre-Partitioning

**Figure 3:** Cost and management burdens when using a specialized OLTP database system.

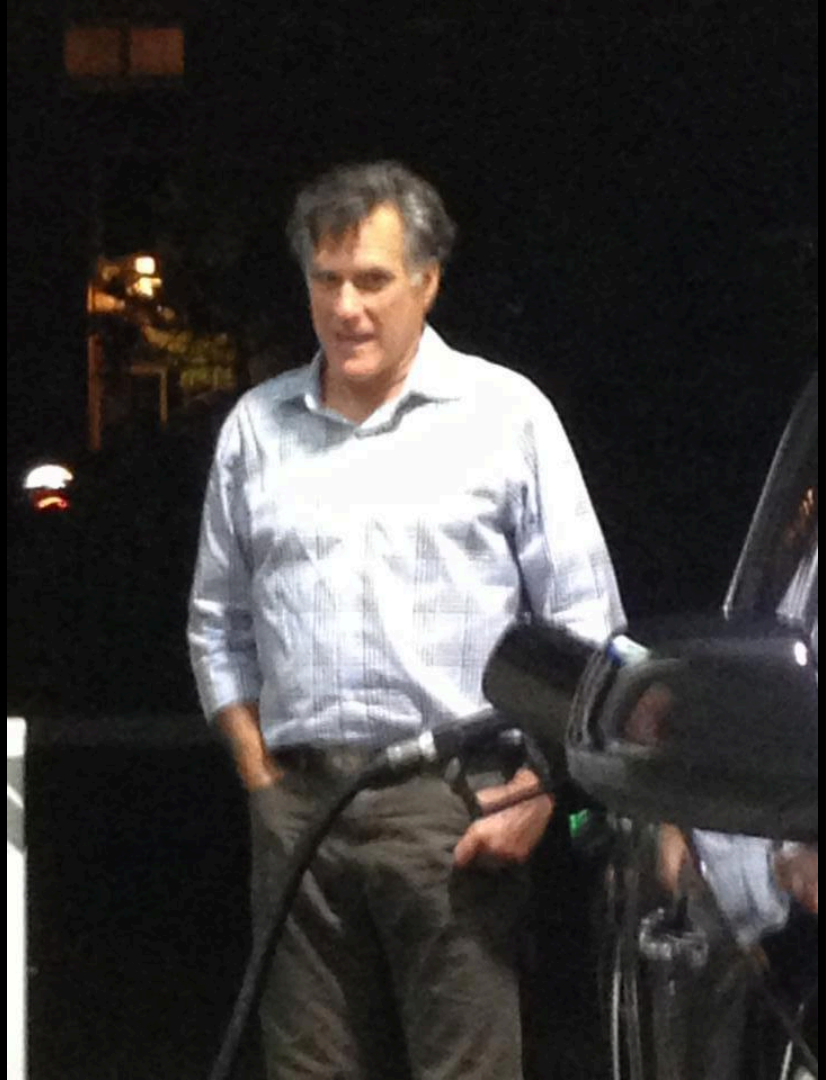

# 99%

Apple
Omniture
foundationdb
Fundings & Exits
Popular Posts

# Apple Acquires Durable Database Company FoundationDB

Posted Mar 24, 2015 by Matthew Panzarino (@panzer)



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.

**(a)** Existing Applications **(b)** Hybrid Workloads **(c)** Autonomous++

**Figure 4:** The design goals of a database management system for the 99%.

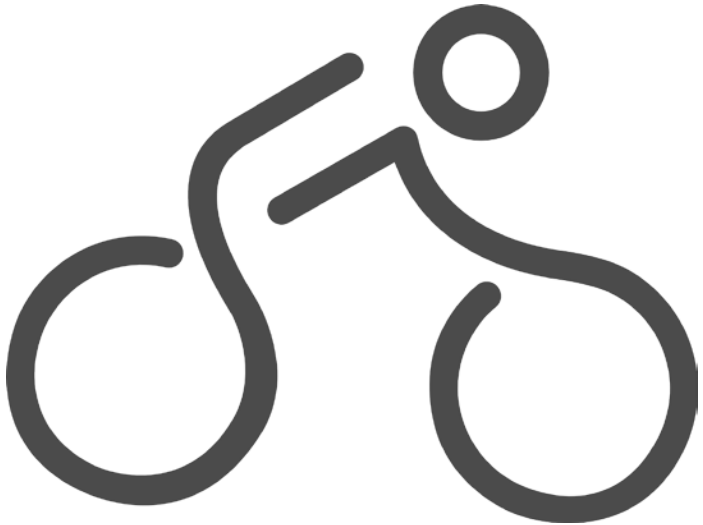

# Peloton



# Data & Execution Model

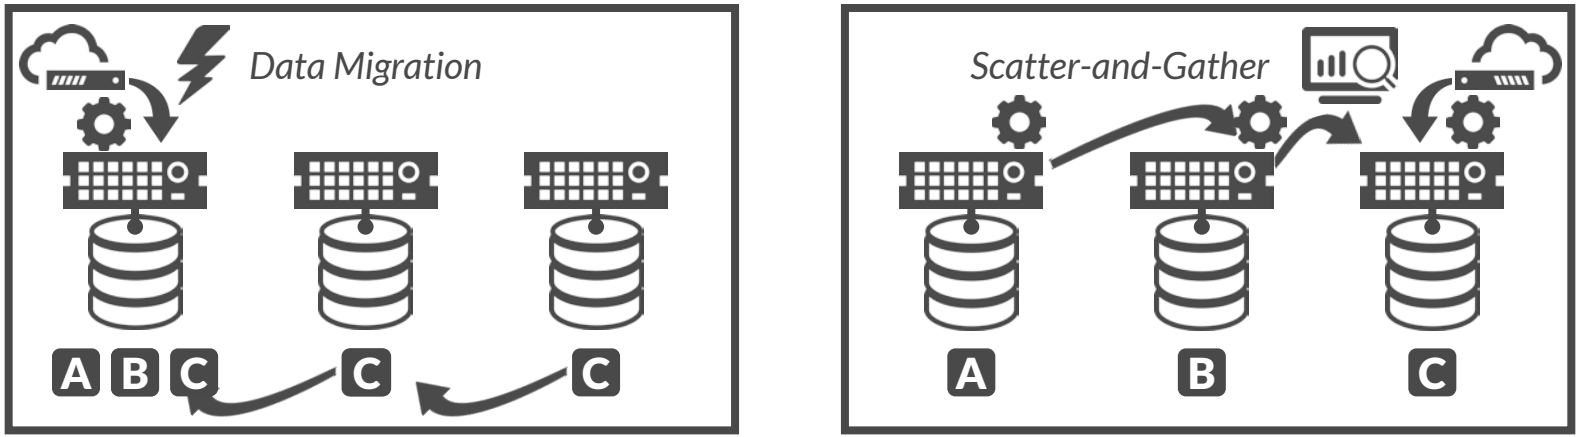


**(a)** OLTP Workload

**(b)** OLAP Workload

**Figure 5:** The database system chooses different data distribution models based on the perceived workload type.

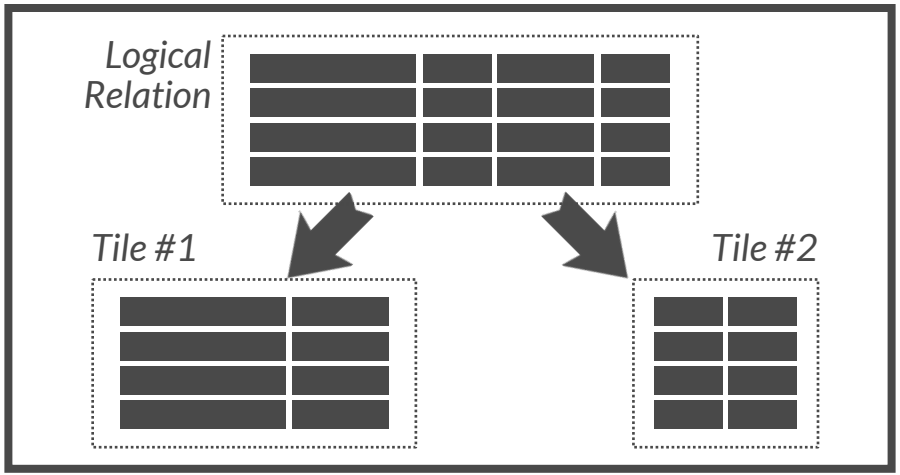


**Figure 6:** Tile-based storage architecture where relations are split into disjoint column blocks.

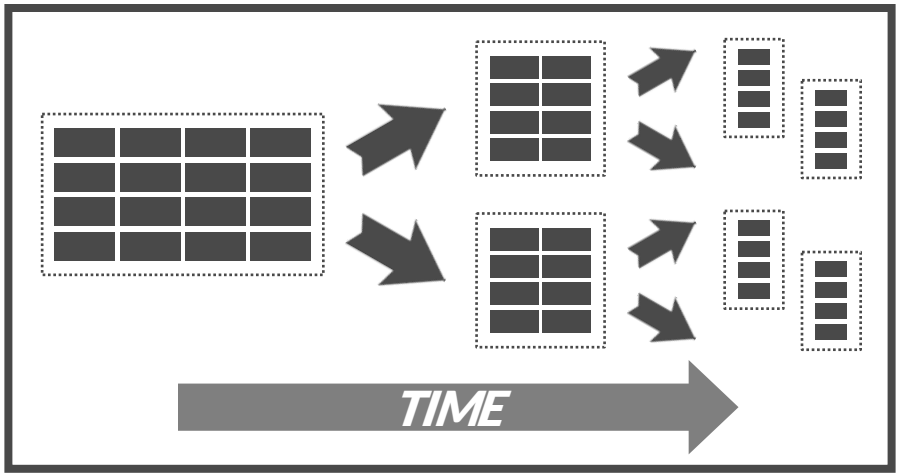


**Figure 7:** Dynamically reorganizing the physical layout of tuples based on the application's access patterns.

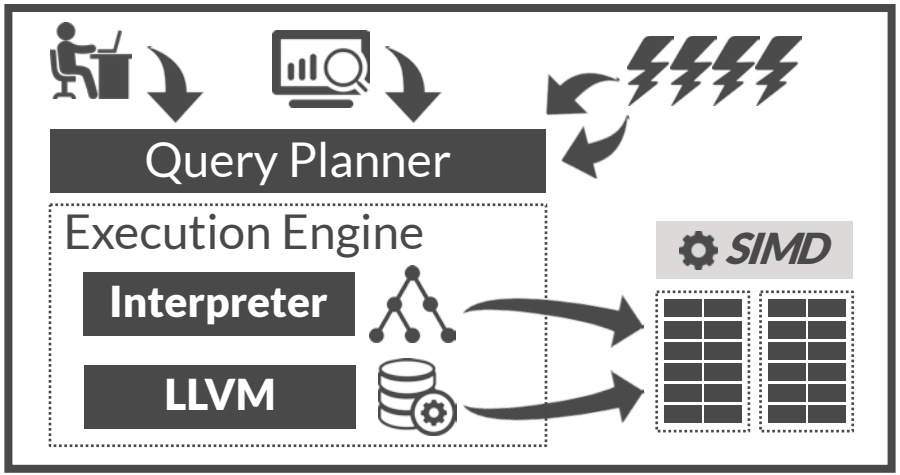


**Figure 8:** Modular query engine that supports interpretation, SIMD execution, and LLVM-based plan compilation.

# Automatic Optimization + Tuning

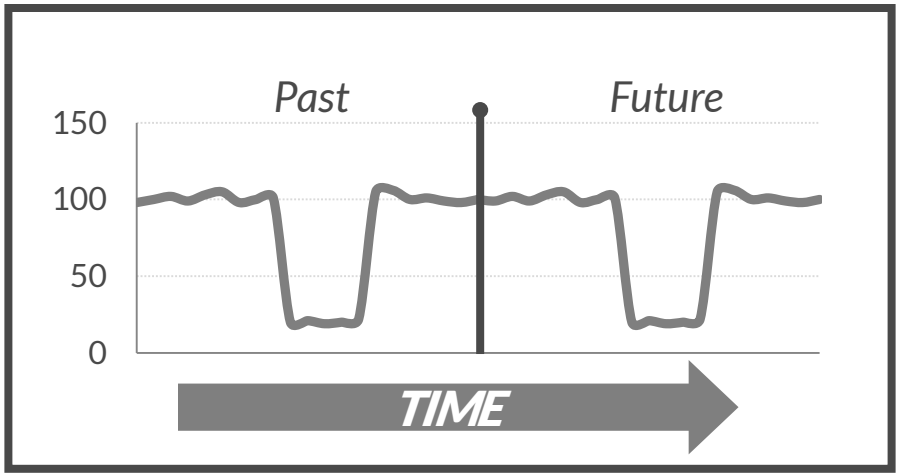


**Figure 9:** An illustration of a database performance metric time-series.

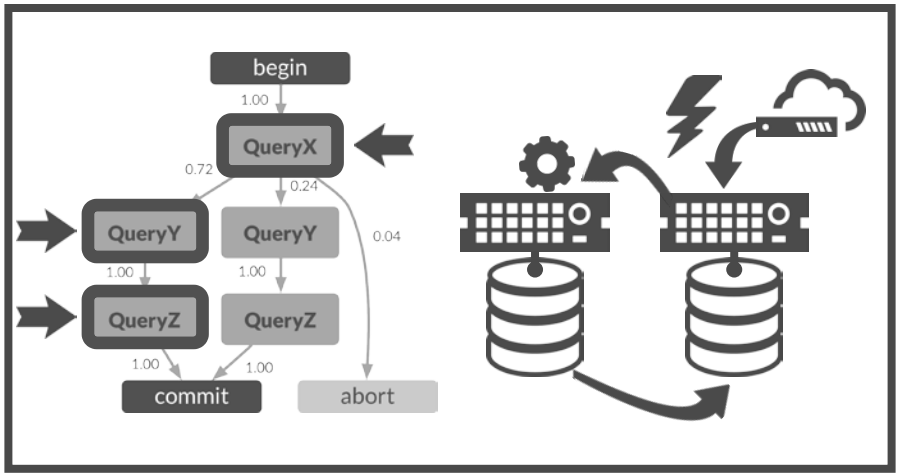


**Figure 10:** Using predictive analytics in OLTP workloads to speculatively execute queries on remote nodes.

*[VLDB 2011]*

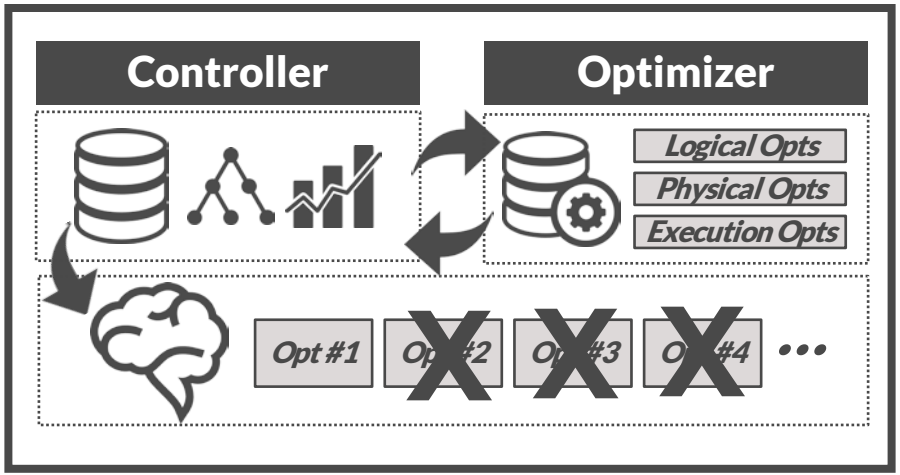


**Figure 11:** A DBMS process controller based on the receding horizon model with scenario-based planning.

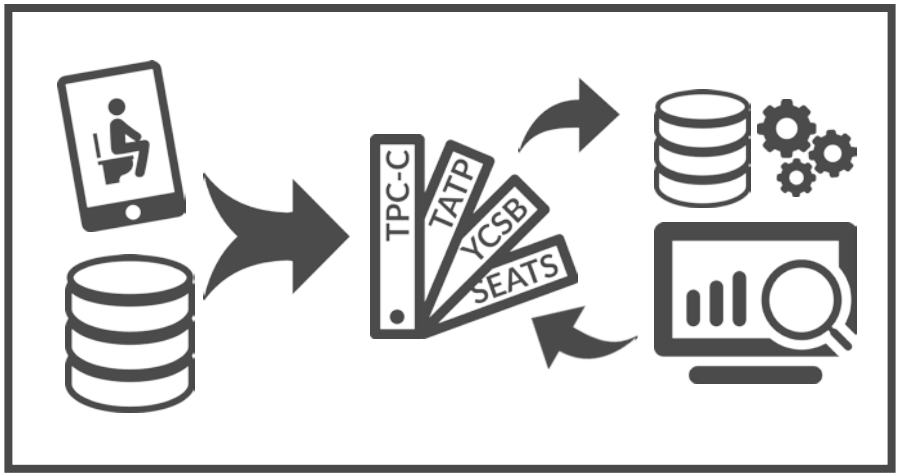


**Figure 12:** Automatic database system configuration tuning using the OLTP-Bench framework.

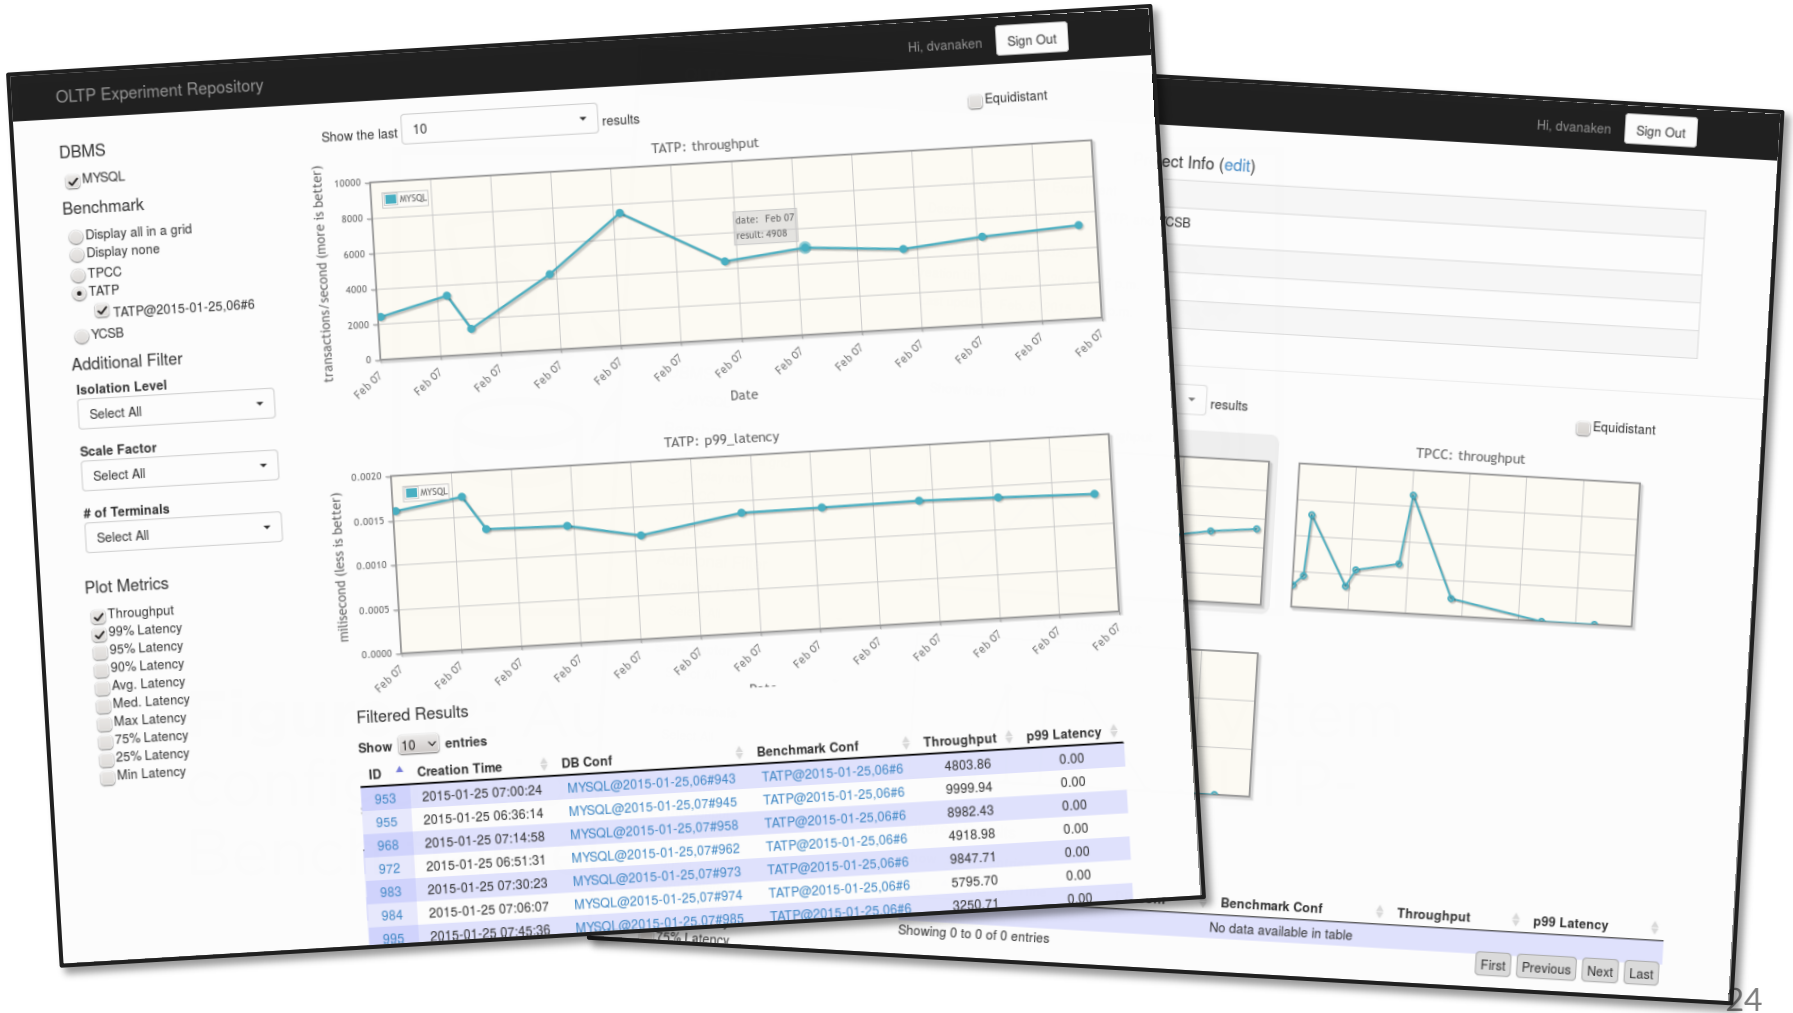
OLTP Experiment Repository
Hi, dvanaken
Sign Out
DBMS
MYSQL
Benchmark
Display all in a grid
Display none
TPCC
TATP
TATP@2015-01-25,06#6
YCSB
Additional Filter
Isolation Level
Select All
Scale Factor
Select All
# of Terminals
Select All
Plot Metrics
Throughput
99% Latency
95% Latency
90% Latency
Avg. Latency
Med. Latency
Max Latency
75% Latency
25% Latency
Min Latency
Show the last 10 results
Equidistant
TATP: throughput
transactions/second (more is better)
MYSQL
date: Feb 07
result: 4908
Date
TATP: p99_latency
millisecond (less is better)
Filtered Results
Show 10 entries
ID
Creation Time
DB Conf
Benchmark Conf
Throughput
p99 Latency
953 2015-01-25 07:00:24 MYSQL@2015-01-25,06#943 TATP@2015-01-25,06#6 4803.86 0.00
955 2015-01-25 06:36:14 MYSQL@2015-01-25,07#945 TATP@2015-01-25,06#6 9999.94 0.00
968 2015-01-25 07:14:58 MYSQL@2015-01-25,07#958 TATP@2015-01-25,06#6 8982.43 0.00
972 2015-01-25 06:51:31 MYSQL@2015-01-25,07#962 TATP@2015-01-25,06#6 4918.98 0.00
983 2015-01-25 07:30:23 MYSQL@2015-01-25,07#973 TATP@2015-01-25,06#6 9847.71 0.00
984 2015-01-25 07:06:07 MYSQL@2015-01-25,07#974 TATP@2015-01-25,06#6 5795.70 0.00
995 2015-01-25 07:45:36 MYSQL@2015-01-25,07#985 TATP@2015-01-25,06#6 3250.71 0.00
Hi, dvanaken
Sign Out
ect Info (edit)
CSB
results
Equidistant
TPCC: throughput
Benchmark Conf
Throughput
p99 Latency
No data available in table
Showing 0 to 0 of 0 entries
First
Previous
Next
Last

# Larger-than-Memory Storage

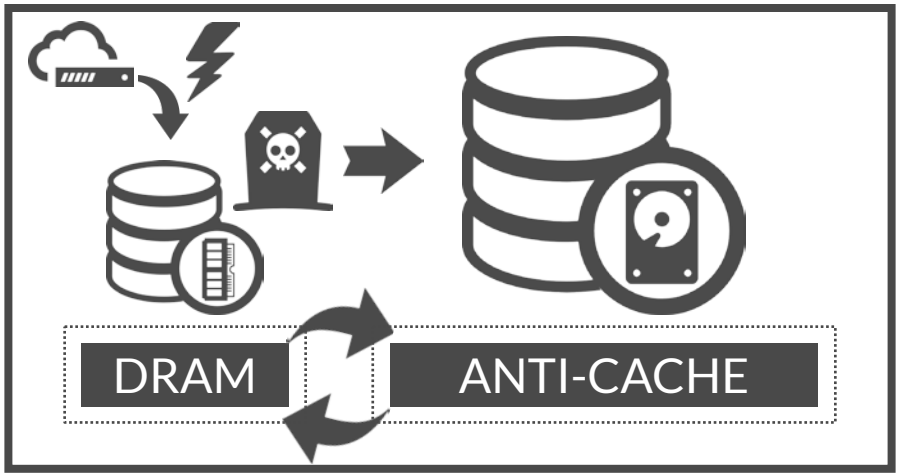


**Figure 13:** New data is first stored in DRAM and then is migrated to the disk-resident anti-cache over time.

*[VLDB 2013]*

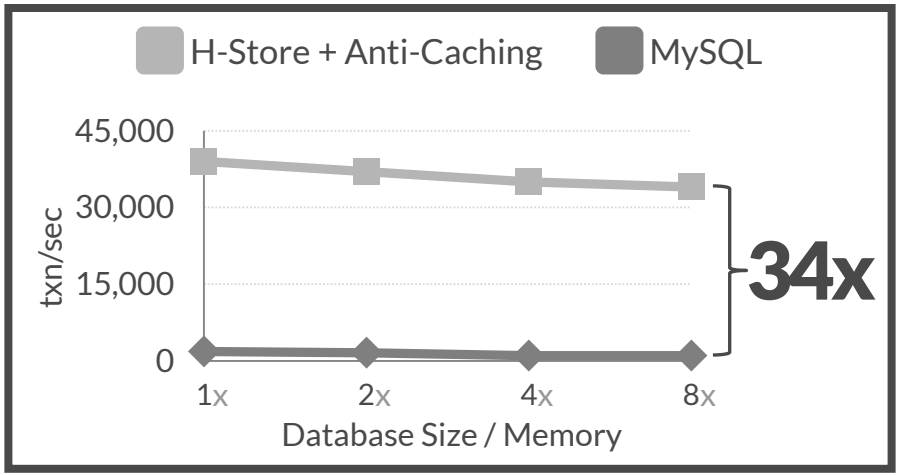


**Figure 14:** TPC-C throughput of H-Store with anti-caching versus MySQL for different database sizes.

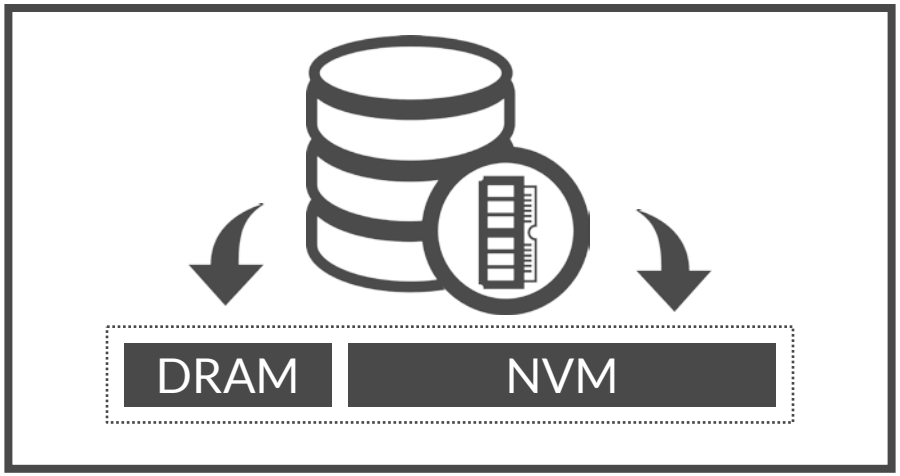


**Figure 15:** Extending the DBMS's address space with NVM. The execution engine treats tuples the same regardless of whether they reside in DRAM or NVM.

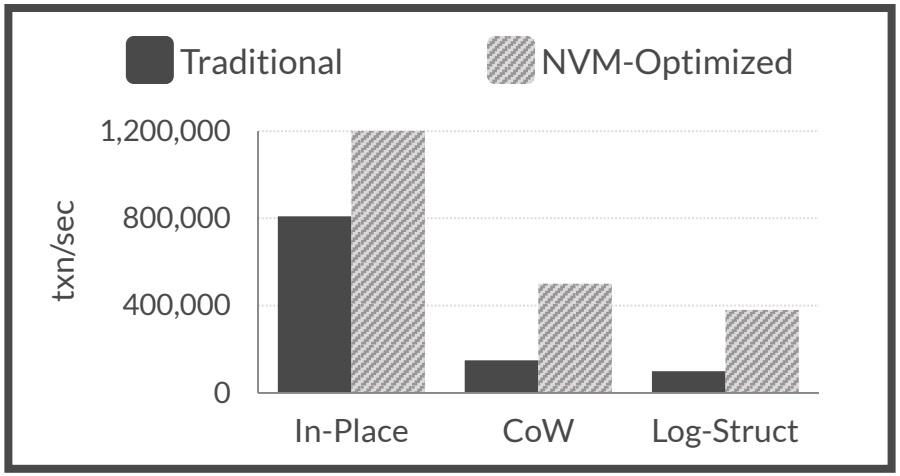


**Figure 16:** Comparison of storage managers executing a YCSB workload. The NVM-optimized engines use byte-addressable persistent data structures.

*[SIGMOD 2015]*

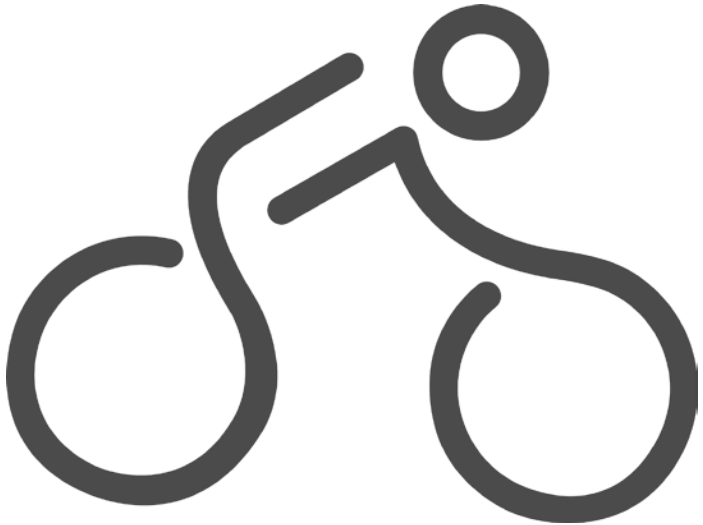

# Peloton

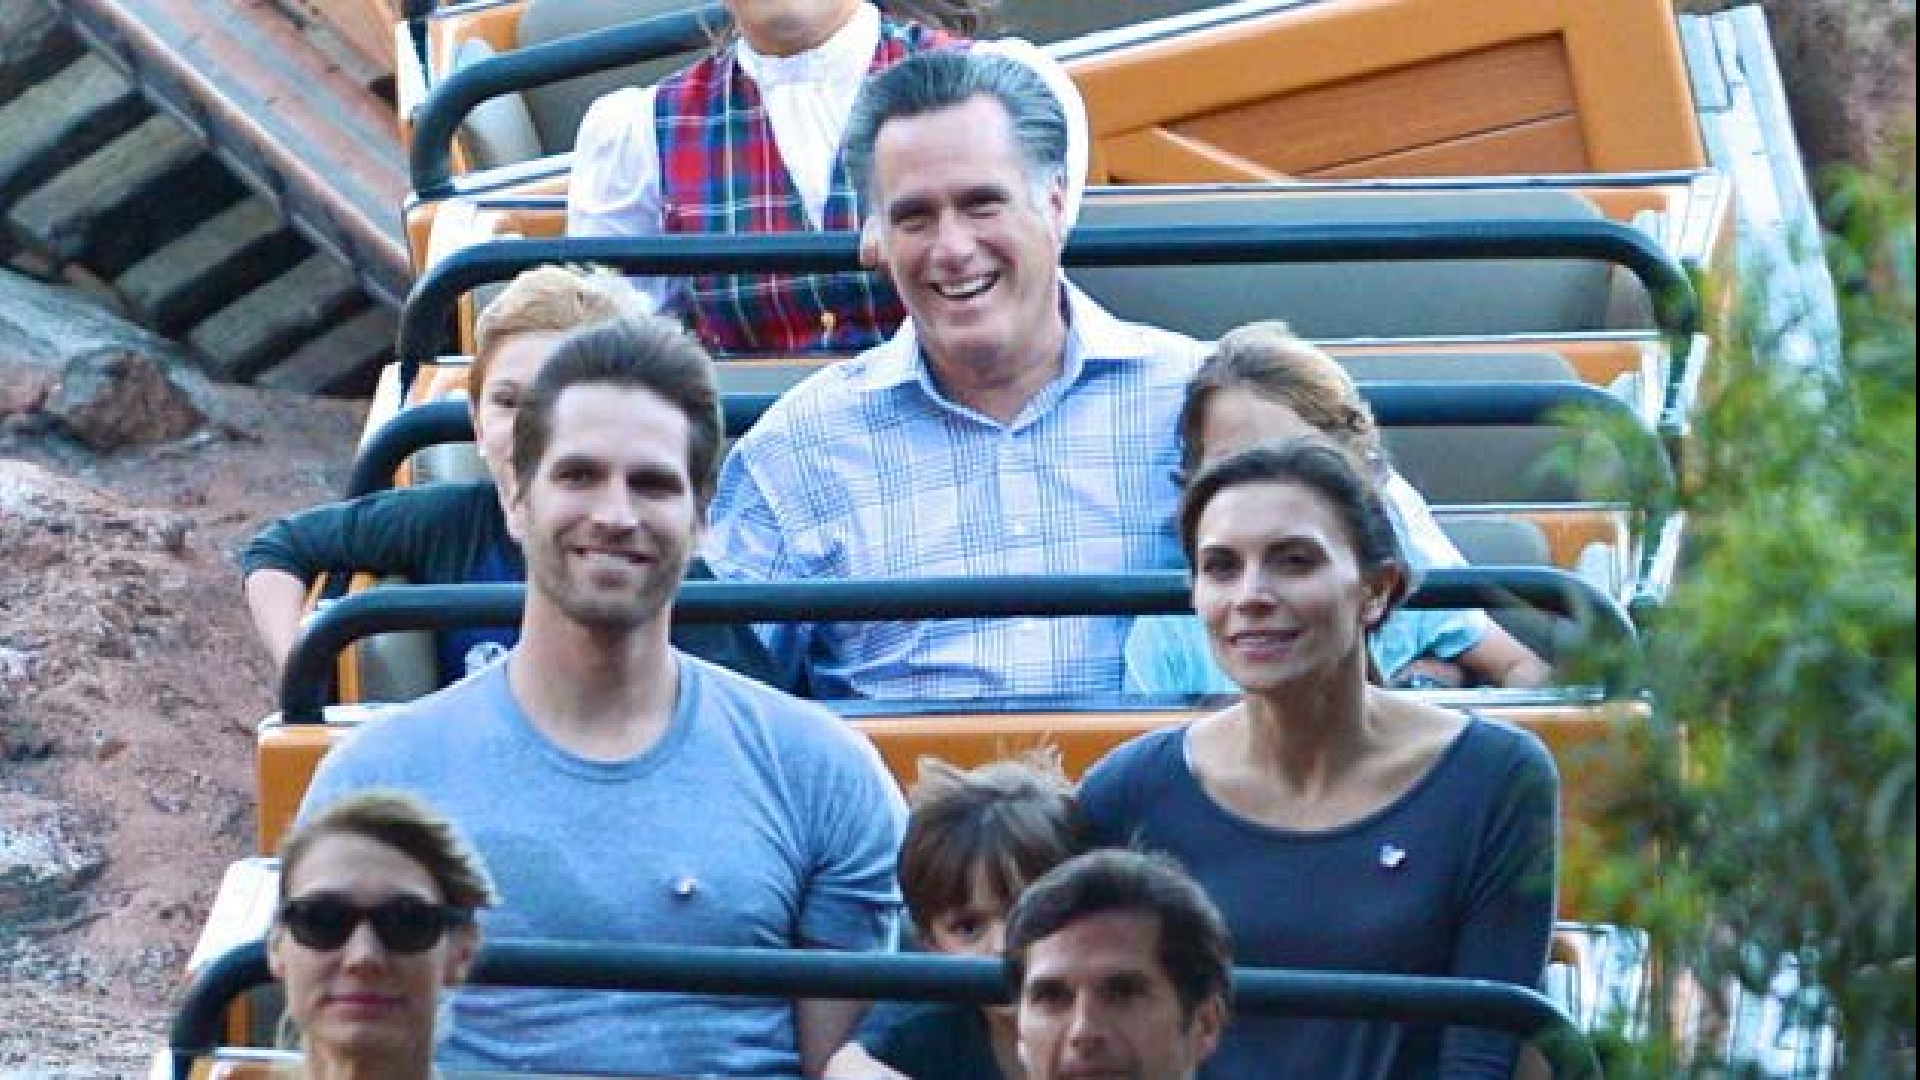

# END

@andy_pavlo