

**Chaitan Baru**, Milind Bhandarkar, Raghunath Nambiar, Meikel Poes, Tilmann Rabl



# **INTRODUCING THE BIG DATA BENCHMARKING COMMUNITY**



# Objectives

- Lay the ground for development of industry standards for measuring the effectiveness of hardware and software technologies dealing with Big Data
- Exploit synergies between benchmarking efforts
- Offer a forum for presenting and debating platforms, workloads, data sets and metrics relevant to Big Data

The banner features the text 'Workshop on Big Data Benchmarking' in a glowing, green-to-blue gradient font. The background is a dark blue with a pattern of binary code and faint, glowing lines. In the top right corner, the text 'WBDB2012' is visible in a lighter blue, semi-transparent font.

# Workshop on Big Data Benchmarking

- First workshop on Big Data Benchmarking, May 8-9, 2012, San Jose
- NSF + Industry sponsorship
  - Seagate, Greenplum, NetApp, Mellanox, Brocade (host)
- Meeting structure
  - 6 x 15mts invited presentations
  - 35 x 5mts lightning talks
  - Afternoon: Structured discussion sessions
- 60 attendees; 45 organizations
- *Setting the Direction for Big Data Benchmark Standards*, Baru, Bhandarkar, Nambiar, Poess, Rabl, TPC Technical Committee Workshop, VLDB 2012.

# WBDB Discussion Topics

- Audience: Who is the audience for this benchmark?
- Application: What application should we model?
- Single benchmark spec: Is it possible to develop a single benchmark to capture characteristics of multiple applications
- Component vs. end-to-end benchmark. Is it possible to factor out a set of benchmark “components”, which can be isolated and plugged into an end-to-end benchmark(s)?
- Paper and Pencil vs Implementation-based. Should the implementation be specification-driven or implementation-driven?
- Reuse. Can we reuse existing benchmarks?
- Benchmark Data. Where do we get the data from?
- Innovation or competition? Should the benchmark be for innovation or competition?

# Workshop Series

- Second workshop, December 17-18, Pune, India
  - Hosted by Persistent Systems Ltd.
  - CFP [clds.ucsd.edu/wbdb2012.in](http://clds.ucsd.edu/wbdb2012.in)
  - Paper submission date: Oct 1, 2012
- Third workshop, late June 2013, Xi'an, China
  - Hosted by the Shanxi Supercomputing Center
  - CFP [clds.ucsd.edu/wbdb2012.cn](http://clds.ucsd.edu/wbdb2012.cn)
- → Big Data Benchmarking Community (BDBC)

# Big Data Benchmarking Community Participants

- Actian
- AMD
- Argonne National Labs
- BMMsoft
- Brocade
- CA Labs
- Cisco
- Cloudera
- CMU
- Convey Computer
- CWI/Monet
- DataStax
- Dell
- EPFL
- Facebook
- GaTech
- Google
- Greenplum
- Hortonworks
- Hewlett-Packard
- IBM
- IndianaU/HTRF
- InfoSizing
- Intel
- Johns Hopkins U.
- LinkedIn
- MapR/Mahout
- Mellanox
- Microsoft
- NetApp
- Netflix
- NIST
- NSF
- OpenSFS
- Oracle
- Ohio State U.
- Paypal
- PNNL
- Red Hat
- San Diego Supercomputer Center
- SAS
- Seagate
- Sholl
- SLAC
- SNIA
- Teradata
- Twitter
- Univ. of Minnesota
- UC Berkeley
- UC Irvine
- UC San Diego
- Univ of Passau
- Univ of Toronto
- Univ. of Washington
- VMware
- WhamCloud
- Yahoo !

# Community Conference Calls

- Conference calls every two weeks. Open to anyone interested
- Information at: [clds.ucsd.edu/bdbc/community](http://clds.ucsd.edu/bdbc/community)
- Contact [tilmann.rabl@utoronto.ca](mailto:tilmann.rabl@utoronto.ca) for scheduling your talk
- Contact Information
  - Call-in number: +1-866-740-1260
  - Access code: 5345802
  - Regular date: Every other Thursday, 2 pm – 3 pm Pacific Coast Time
- Upcoming presentations:
  - Sep 20, Y. Chen, F. Raab, From TPC-C to Big Data Benchmarks
  - Oct 04, F. Stegmaier, The Code Project

# Thank You

- Visit our Poster...
- BDBC Steering Committee:
  - **Chaitan Baru**, CLDS, UCSD, [baru@sdsc.edu](mailto:baru@sdsc.edu)
  - **Milind Bhandarkar**, Greenplum, EMC, [milind.bhandarkar@emc.com](mailto:milind.bhandarkar@emc.com)
  - **Raghunath Nambiar**, Cisco Systems Inc., [rnambiar@cisco.com](mailto:rnambiar@cisco.com)
  - **Meikel Poess**, Oracle Corporation, [meikel.poess@oracle.com](mailto:meikel.poess@oracle.com)
  - **Tilmann Rabl**, MSRГ, University of Toronto, [tilmann.rabl@utoronto.ca](mailto:tilmann.rabl@utoronto.ca)
- More info: [clds.sdsc.edu/bdbc](http://clds.sdsc.edu/bdbc)







# Audience: Who is the primary audience for a big data benchmark?

- Customers
  - ➔ Workload should preferably be expressed in
    - English
    - Or, a declarative Language (unsophisticated user)
    - But, not a procedural language (sophisticated user)
  - Want to compare among different vendors
- Vendors
  - Would like to sell machines/systems based on benchmarks
- Computer science/hardware research is also an audience
  - Niche players and technologies will emerge out of academia
  - Will be useful to train students on specific benchmarking

# Applications: What application should we model?

- Possibilities


- An application that somebody could donate
- An application based on empirical data
  - Examples from scientific applications
- Multi-channel retailer-based application, like the amended TPC-DS for Big Data?
  - Mature schema, large scale data generator, execution rules, audit process exists.
- “Abstraction” of an Internet-scale application, e.g. data management at the Facebook site, with synthetic data


# Single Benchmark vs Multiple

- Is it possible to develop a single benchmark to represent multiple applications?
- Yes, but not desired if there is no synergy between the benchmarks, e.g. say, at the data model level
  - Synthetic Facebook application might provide context for a single benchmark
  - Click streams, data sorting/indexing, weblog processing, graph traversals, image/video data, ...




# Component benchmark vs. end-to-end benchmark

- Are there components that can be isolated and plugged into an end-to-end benchmark?
  - The benchmark should consist of individual components that ultimately make up an end-to-end benchmark
  - The benchmark should include a component that extracts large data
    - Many data science applications extract large data and then visualize the output
    - Opportunity for “pushing down” viz into the data management system
- 



# Paper and Pencil / Specification driven versus Implementation driven

- Start with an implementation and develop specification at the same time
  - Some post-workshop activity has begun in this area
    - Data generation; sorting; some processing
- 

# Where Do we Get the Data From?

- Downloading data is not an option
- Data needs to be generated (quickly)
- Examples of actual datasets from scientific applications
  - Observational data (e.g. LSST), simulation outputs
- Using existing data generators (TPC-DS, TPC-H)
- Data that is generic enough with good characteristics is better than specific data

# Should the benchmark be for innovation or competition?

- Innovation and competition are not mutually exclusive
  - Should be used for both
  - The benchmark should be designed for competition, such a benchmark will then also be used internally for innovation
- TPC-H is a prime example of a benchmark model that could drive competition and innovation (if combined correctly)



# Can we reuse existing benchmarks?

- Yes, we could but we need to discuss:
  - How much augmentation is necessary?
  - Can the benchmark data be scaled
  - If the benchmark uses SQL, we should not require it
- Examples: but none of the following could be used unmodified
  - Statistical Workload Injector for MapReduce (SWIM)
  - GridMix3 (lots of shortcomings)
    - Open source
  - TPC-DS
  - YCSB++ (lots of shortcomings)
  - Terasort – strong sentiment for using as an “end-to-end” scenario

**Big Data Benchmark Data Model**

- TPC-DS
  - Decision support benchmark from Transaction Processing Performance Council
  - [http://www.tpc.org/tpcds/spec/tpcds\\_1.1.0.pdf](http://www.tpc.org/tpcds/spec/tpcds_1.1.0.pdf)
- Why build on top of TPC-DS?
- Volume :
  - No theoretical limit
  - Tested up to 100 TB
- Velocity : rolling updates
- Variety
  - Rich relational model and data tables, text tables, ...
  - Easy to add the other two sources

4

TERADATA | aster data

# Keep in mind principles for good benchmark design

- Self-scaling, e.g. TPC-C
- Comparability between scale factors
  - Results should be comparable at different scales
- Technology agnostic (if meaningful to the application)
- Simple to run

## TPC

+ Longevity: TPC-C has carried the load for 20 years

+ Comparability

- Audit requirements and strict detailed run rules mean one can compare results published by two different entities


+ Scaling

- Results just as meaningful at the high-end of the market as at the low-end; as relevant on clusters as on single servers

- Hard and expensive to run

- No kit

- DeWitt clauses



# Extrapolating Results

- TPC benchmarks typically run on “over-specified” systems
  - i.e. Customer installations may have less hardware than benchmark installation (SUT)
- Big Data Benchmarking may be opposite
  - May need to run benchmark on systems that are smaller than customer installations
    - Can we extrapolate?
  - Scaling may be “piecewise linear”. Need to find those inflexion points