



Lessons from Industry Standard Benchmarking

Raghunath Nambiar
Distinguished Engineer, Data Center Group, Cisco Systems, Inc

December 17-18, 2012 in Pune, India

**SECOND WORKSHOP ON
BIG DATA BENCHMARKING**

Center for Large-Scale Data Systems [CLDS]
San Diego Supercomputer Center, UC San Diego

Agenda

- Industry Standard Benchmarks
- A closer look the TPC
- Changing technology landscape
- Key considerations for Big Data benchmarks

Benchmarks

- Demands for them have existed since buyers were first confronted with the choice between purchasing one system over another

Historically we have seen that industry standard benchmarks enable healthy competition that results in product improvements and the evolution of brand new technologies.

Categories

- Industry standard benchmarks
 - Consortia driven development (Democratic process in decision making), Verifiable, Audit process, Product agnostic, Enables cross technology comparison

TPC Transaction Processing
Performance Council



Storage Performance Council

- Application benchmarks
 - Vendor driven, vertical focused
 - Examples: VMMark, SAP Standard Application Benchmarks, Oracle Applications Benchmarks
- Synthetic Workloads
 - Open source based tools, widely adapted by engineering community
 - Examples: IOMeter, Netperf, Terasort

Viewpoints

- Vendor point of view
 - Define the level playing field for competitive analysis
 - Monitor release to release progress
 - Product developments and enhancements
- Customer point of view
 - Cross-vendor comparisons (performance, cost, power)
 - Evaluate new technologies
- Researcher point of view
 - Known, measurable and repeatable workloads
 - Optimizations can impact products

Benchmarks and Technology Landscape

Driven by technology and industry demands

1985: Debit Credit Benchmark

1988: TPC and SPEC formed

1990s (Client server, Internet)

- Processor
- Transaction Processing
- File server, Web server
- Data Warehouse


2000s (Web 2.0, Virtualization)

- Energy Efficiency
- Virtualization
- Complex Systems
- Cloud

2010s (Connected world, Analytics)

- Big Data
- Internet of things
- Software defined “everything”

Worldwide Cost to Power and Cool Server Installed Base, 1996–2010



Source: IDC, 2007

Did you know the number of virtual machine shipments now exceeds the number of physical server shipments ?

TPC Membership

Full Members

Associate Members

				
---	---	--	---	---


- Database centric
- Represented by major systems and database companies
- Independent audit process prior to publication
- Mandatory performance, price-performance metric
- Optional energy efficiency and virtualization metrics

TPC Timeline

Benchmark Standards																									
TPC-A			Active	Active	Active	Active	Active	Active																	
TPC-B			Active	Active	Active	Active	Active																		
TPC-C					Active	Active	Active	Active	Active	Active	Active	Active	Active	Active	Active	Active	Active								
TPC-D							Active	Active	Active	Active	Active														
TPC-R										Active	Active	Active	Active	Active											
TPC-H										Active	Active	Active	Active	Active	Active	Active	Active								
TPC-W										Active	Active	Active	Active												
TPC-App														Active	Active	Active									
TPC-E														Active	Active	Active	Active								
TPC-DS																	Active								
TPC-VMS																	Active								
Common Specifications																									
Pricing																	Active								
Energy																	Active								
Developments in Progress																									
TPC-DI																	In Progress								
TPC-VMC																	In Progress								
TPC-V																	In Progress								
	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012

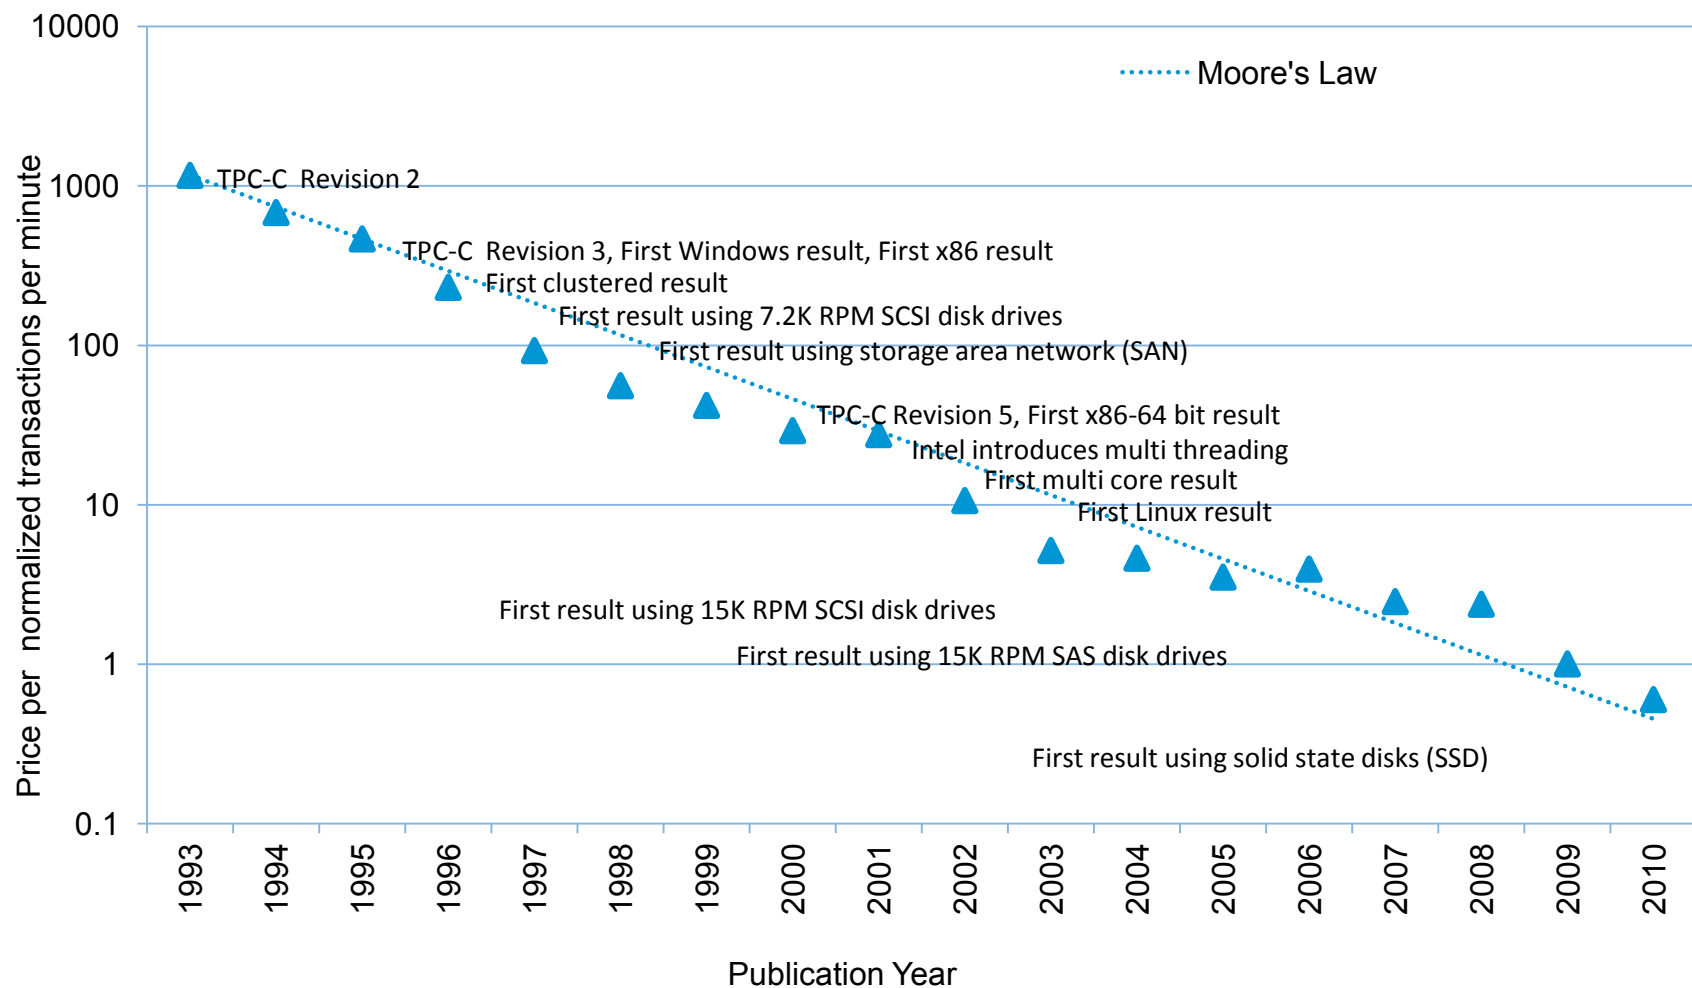
Obsolete
Active
Common Specifications
In Progress

TPC-C Performance vs. Moore's Law



Reference: R. Nambiar, M. Poess, Transaction Performance vs. Moore's Law: A Trend Analysis: <http://www.springerlink.com/content/fq6n225425151344/>

TPC-C Price-Performance vs. Moore's Law



Reference: R. Nambiar, M. Poess, Transaction Performance vs. Moore's Law: A Trend Analysis: <http://www.springerlink.com/content/fq6n225425151344/>

SPEC Membership

SPEC Members:

*Acer Inc. * Action S.A. * Advanced Micro Devices * Apple Inc. * ASUSTeK Computer Inc. * Avere Systems * Bull S.A. * China Telecom Corporation Limited (Beijing Research Institute) * Cisco Systems, Inc. * Dell, Inc. * E4 Computer Engineering SPA * EMC * Fujitsu * Hitachi Data Systems * Hitachi Ltd. * HP * Huawei Technologies Co. Ltd. * IBM * IDEAS International * Intel * Itautec S/A * Lenovo * Microsoft * NEC - Japan * NetApp * NVIDIA * Oracle * Platform Computing Inc. * Principled Technologies * The Portland Group * Red Hat * Samsung * SAP AG * SGI * Super Micro Computer, Inc. * SUSE * Symantec Corporation * Unisys * Via Technologies * VMware * Voxeo Corporation **

SPEC Associates:

*Academia Sinica, Institute of Information Science * Austrian Energy Agency * Argonne National Laboratory * Charles University * Citrix * Department of Veteran's Affairs - Corporate Data Center Operations * Dresden University of Technology ZIH * Indiana University * Institute for Information Industry Taiwan * JAIST * Karlsruhe Institute of Technology * Leibniz Rechenzentrum - Germany * National University of Singapore * Ohio State University * Purdue University * Technische Universität Darmstadt * University of Aizu - Japan * University of California - Berkeley * University of Cologne * University of Illinois at Urbana-Champaign * University of Maryland * University of Miami * University of Pavia * University of Texas at Austin * University of Tsukuba * University of Virginia * Virginia Polytechnic Institute and State University * Zhejiang University **

SPEC Research Group:

*AMD * Charles University * Cisco Systems * Dell * Ghent University * Hasso Plattner Institute * Hewlett-Packard * IBM * Imperial College London * Indian Institute of Technology, Bombay * Institute for Information Industry, Taiwan * Intel * Karlsruhe Institute of Technology * McGill University * NICTA * Oracle * Purdue University * Red Hat * Salesforce.com * San Francisco State University * SAP * Siemens Corporation * Technische Universität Darmstadt * University of Alberta * University of California at Berkeley * University of Coimbra * University of Kiel * University of Mannheim * University of Minnesota * University of Pavia * University of Texas at Austin * VMware * Zhejiang University **

System centric

Broad representation from industry and research


Peer audit process prior to publication

Primary metric is: performance

Separate benchmark for energy efficiency and virtualization

Technology Landscape is Changing

the winner of Time's man
of the year award for
1982 ...




the winner of Time's man
of the year award for
2006 ...



Internet Users in the World

Distribution by World Regions - 2012 Q2



Source: Internet World Stats - www.internetworldstats.com/stats.htm

Basis: 2,405,518,376 Internet users on June 30, 2012

Copyright © 2012, Miniwatts Marketing Group

$2,405,518,376 / 7,017,846,922 = 34.3 \%$


it took ...

53 years to sell 100 million radios

14 years to sell 100 million TVs

but .. in 15 years ...

there are 15 billion devices connected to the Internet



that's 2.2 devices for every man, woman,
and child on the planet earth

if



were a country ...



1. China (1.339 billion)
2. India (1.218 billion)
3. Facebook (900 million) 
4. United States (311 million)
5. Indonesia (237 million)
6. Brazil (190 billion)
7. Pakistan (175 million)
8. Nigeria (158 million)
9. Bangladesh (150 million)
10. Russia (142 million)



Processing Movie AVATAR

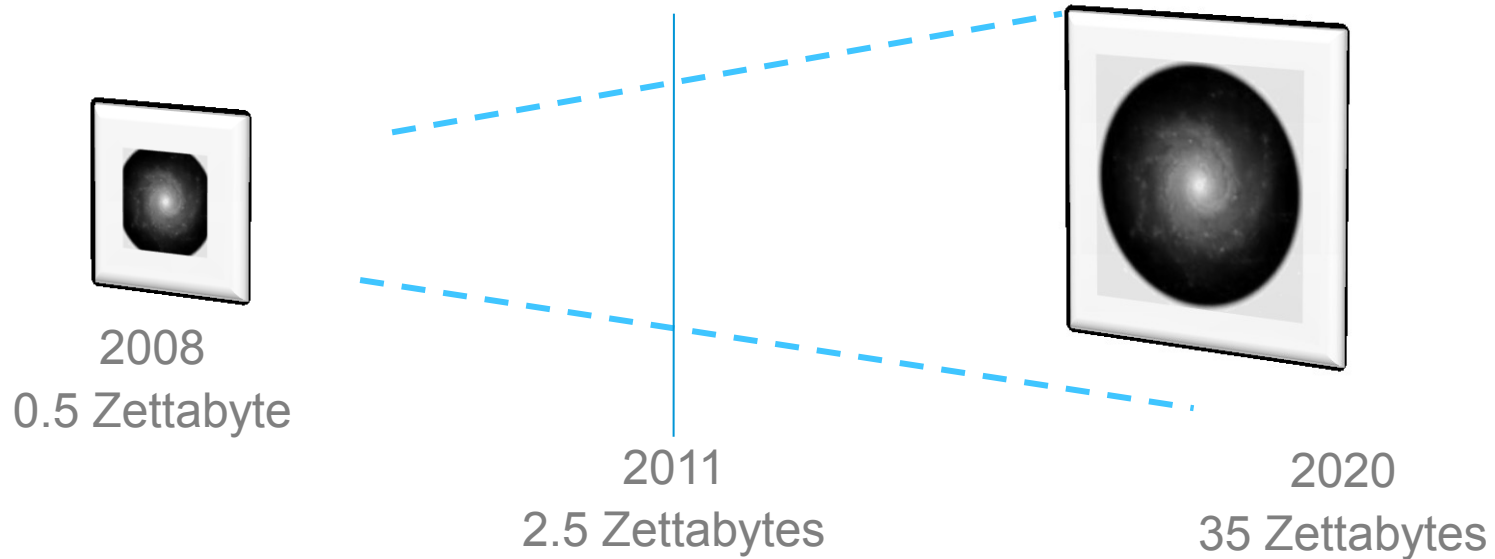
4352 servers
34,816 processor-cores
102 terabytes of RAM
3 Petabytes of storage



Comprehensive statistics on more than
250 countries and entities

Unknown number of classified information

How Big is the Digital Universe ?




1 Zettabyte = 1 099 511 627 776 Gigabytes
= 1 Billion **1TB** Disk Drives


How many disk drives were sold in 2011 ?

Global IP Traffic

Global IP Traffic



Per Capita Internet Traffic



In 2016, equivalent of all movies ever made will cross global IP networks every 3 minutes

Source: Cisco

Benchmark Acceleration Initiatives

- TPC

Technology conference initiative on performance evaluation and benchmarking

2009 (Lyon), 2010 (Singapore), 2011 (Seattle), 2012 (Istanbul)
2013 (Trento) – planned

- SPEC

- SPEC Research
- ICPE

- WBDB

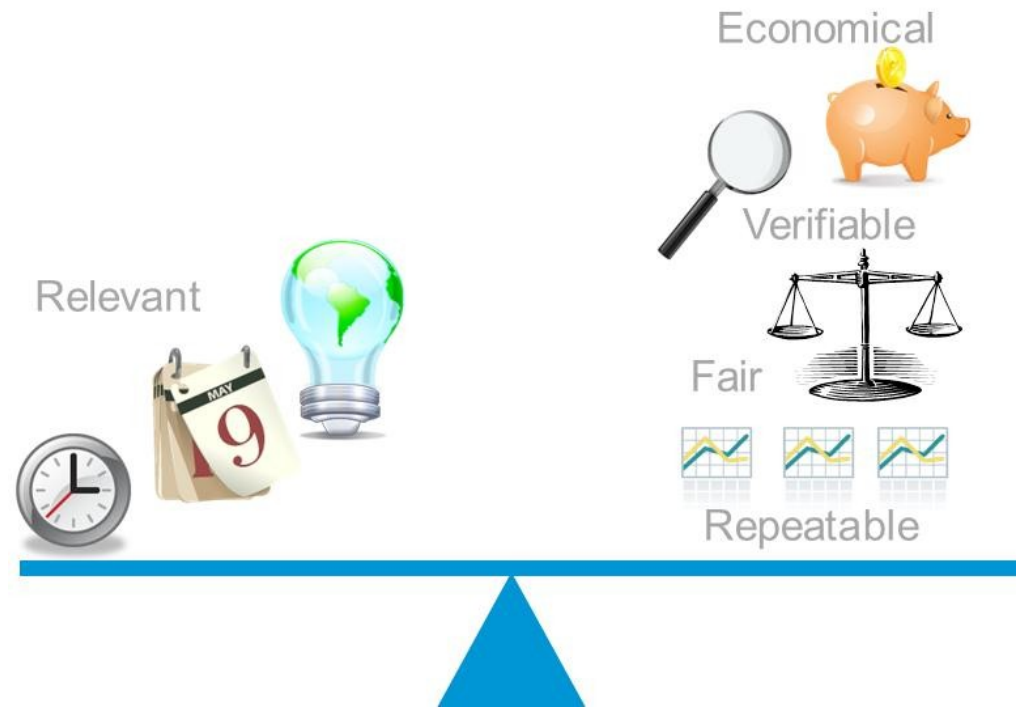
- First important step towards the development of a set of benchmarks for providing objective measures of the effectiveness of hardware and software systems dealing with big data applications.

What is Important ?

- Performance
- Cost of ownership
- Energy efficiency
- Floor space efficiency
- Manageability
- User experience

Successful Benchmark Requirements

- Relevant
- Repeatable
- Understandable
- Fair
- Verifiable
- Economical



Reference: K. Huppler, The Art of Building a Good Benchmark, Performance Evaluation and Benchmarking, [LNCS vol. 5895](#), Springer 2009

WBDB

- WBDB 2012

<http://clds.ucsd.edu/wbdb2012>

Chaitanya Baru, Milind Bhandarkar, Raghunath Nambiar, Meikel Poess, Tilmann Rabl: Setting the Direction for Big Data Benchmark Standards, TPCTC 2012, LNCS Vol. 7755, Springer 2012

- WBDB 2012.in

<http://blogs.cisco.com/datacenter/wbdb2012-in/>

- WBDB 2013

<http://clds.ucsd.edu/wbdb2013.cn>

- Bigdata 100 List

Thank you.

