



Exadata Evaluation using Key Performance Metrics

Manfred Drozd
DOAG Jahrestagung 2011
November 2011

About us



Strong cooperation in Exadata projects

- **Benchmark AG**
 - independent company
 - services and products for benchmarking and for performance analysis
 - Benchmark tools to compare performance capabilities of Oracle platforms, e.g. Exadata versus conventional platforms
- **Tradeware AG**
 - Focus on Exa-stack
 - Exadata Competence Center in Switzerland
 - **EMEA Exadata Partner of the Year 2011**

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 The Benchware approach: Key Performance Metrics
- 5 Benchmark Example: conventional platform versus Exadata
- 6 Conclusions

Issues with Proof-of-Concepts



- Time Consuming – Complicated – Expensive
 - Data masking
 - Large user populations
 - Interfaces (SOA, ESB)
 - Consolidation of database servers
- Requires scalable application
- Application may shift behavior on much faster platform, e.g. when moving to Exadata

- A PoC is still NOT reality, but some kind of modeling reality

Proof-of-Concept?




Engine	
Cylinders/valves	4/4
Capacity in ccm	1,995
Stroke/bore in mm	90.0/84.0
Max. output in kW (hp) at 1/min	105 (143)/6,000
Max. torque in Nm at 1/min	190/4,250
Power-to-weight ratio (EU) in kg/hp	10.5

Performance	
Drag (cw)	0.32
Top speed (km/h)	210
Acceleration 0 - 100 km/h (in s)	9.3
Acceleration 0 - 1,000 m (in s)	30.6
Acceleration 80 - 120 km/h in 4th/5th gear (in s)	9.6/12.5



Source: www.bmw.de



copyright © 2011 by benchware.ch

slide 5

No – Instead Key Performance Metrics!





Engine	
Cylinders/valves	4/4
Capacity in ccm	1,995
Stroke/bore in mm	90.0/84.0
Max. output in kW (hp) at 1/min	105 (143)/6,000
Max. torque in Nm at 1/min	190/4,250
Power-to-weight ratio (EU) in kg/hp	10.5

Performance	
Drag (cw)	0.32
Top speed (km/h)	210
Acceleration 0 - 100 km/h (in s)	9.3
Acceleration 0 - 1,000 m (in s)	30.6
Acceleration 80 - 120 km/h in 4th/5th gear (in s)	9.6/12.5



Source: www.bmw.de



copyright © 2011 by benchware.ch


slide 6

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 The Benchware approach: Key Performance Metrics
- 5 Benchmark Example: conventional platform versus Exadata
- 6 Conclusions

Complexity of database platforms



Complexity of database platform

IP-based Network

Bandwidth, latency during remote database mirroring (sync, async) due to switches and sql*net and tcp/ip stack (frame size, ...).

Oracle Database

Different versions, patches and options, about hundred configuration parameters.

FC-based or IB-based Network

Bandwidth, latency during remote storage mirroring (sync, async) due to switches, hubs and distance.

Volume & File Management

Different volume managers (VxVM, ASM) and file Systems (UFS, VxFS, ext3, JFS, ZFS, raw devices), different I/O methods (async, direct), a lot of config parameters (#LLUNS, queue depth, max i/o unit), software striping and/or mirroring, multipathing.


Server & Operating System

Different server Systems, processors and CPU architectures, (x86, IA-64, UltraSparc, SPARC64, Power), #cores, multithreading, main memory, bus architecture. Different operating Systems and patches, over hundred configuration parameters, virtualization of resources.

Storage System

Different storage Systems, storage tiers and storage technology: spindle count and speed, RAID management, cache management, server interface technology, storage system options like remote copy, hardware striping and/or mirroring, virtualization of resources.

Complexity of database platforms



copyright © 2011 by benchware.ch

slide 9

Complexity of database platforms



- Due to the complexity of platforms, the performance of database platforms is not predictable
- Benchmarking is the only way to predict the performance of a database platform

copyright © 2011 by benchware.ch

slide 10

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 The Benchware approach: Key Performance Metrics
- 5 Benchmark Example: conventional platform versus Exadata
- 6 Conclusions

Useless TPC benchmarks



- Strict rules
 - Sizing
 - Usable features
- Unrealistic hardware configuration
- Only one load profile
- Key performance metrics
 - Very synthetic – difficult to understand
 - TPC-H: QphH@Size (Query per hour at given database size)

*"In short, TPC has become vendor-dominated, and it is time for TPC to reinvent itself to serve its customer community",
Michel Stonebraker, Keynote at TPC Technology Conference 2009*

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 **The Benchware approach: Key Performance Metrics**
- 5 Benchmark Example: conventional platform versus Exadata
- 6 Conclusions

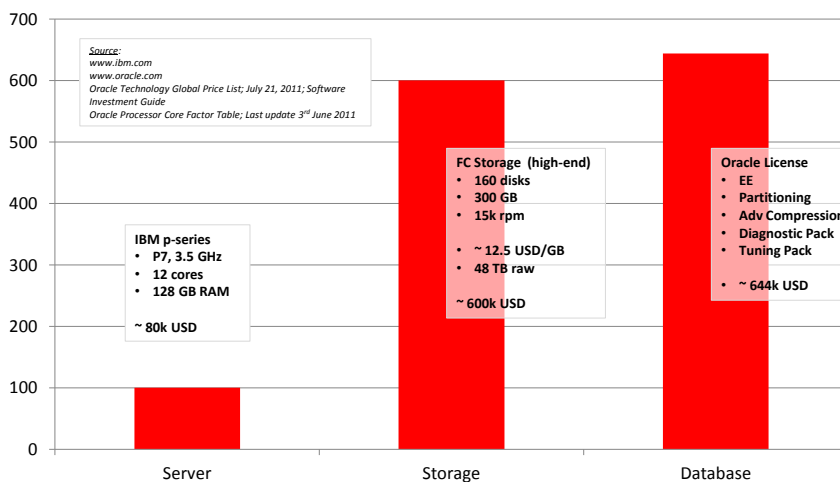
copyright © 2011 by benchware.ch

slide 13

Oracle Platform Investment



Main components: CPU and server, storage, database



copyright © 2011 by benchware.ch

slide 14

The Benchware approach



Library of Oracle benchmark tests, implemented in PL/SQL, Java and SQL

Server Performance CPU-bound Oracle operations	OLTP systems	DWH systems	Metric	Unit
• numeric and string operations in PL/SQL and Java	★★	★	speed	[ns], [s]
• algorithms in PL/SQL and Java	★★	★★	throughput	[ops], [rps]
• in-memory (SGA) SQL operations	★★★	★		
Storage Performance I/O-bound Oracle operations				
• sequential access pattern	★★	★★★	throughput	[mbps], [iops]
• random access pattern	★★★	★	service time	[ms]
Database Performance Mixed resource usage: CPU, memory, storage				
• data load (conventional, bulk)	★★	★★★	throughput	[rps], [tps], [qpm]
• data scan	★	★★★		
• data aggregation, cubes, reports	★★	★★★	service time	[ms], [s]
• OLTP transactions	★★★	★		

[s] seconds
[ms] milli seconds (10⁻³)
[ns] nano seconds (10⁻⁹)

[rps] rows per second
[tps] transactions per second
[qpm] queries per minute

★ less important
★★ important
★★★ very important

copyright © 2011 by benchware.ch

slide 15

The Benchware approach



Following customers have used our approach

IT vendors

- HP European Benchmark Center
- Local swiss organization of HDS, IBM, Intel, Sun Microsystems

Banking

- Julius Bär, Pictet, Sarasin, Vontobel, Migros¹⁾, Raiffeisen, UBS, Credit Suisse, ING-DiBa

Public sector:

- Swiss Post, Cineca Italy¹⁾, Kanton Basle

¹⁾ Exadata projects

High-Tech

- Nagravision/Kudelski¹⁾

Insurances

- Concordia, CSS, Mobiliar

Chemistry

- EMS

Telecommunication

- Swisscom¹⁾, Sunrise

Trading

- Charles Vögele, Fenaco, Koch Neff Volckmar Germany

copyright © 2011 by benchware.ch

slide 16

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 The Benchware approach: Key Performance Metrics
- 5 **Benchmark Example: conventional platform versus Exadata**
- 6 Conclusions


copyright © 2011 by benchware.ch

slide 17

Benchmark Results



CPU and Server: PL/SQL arithmetic mix data type NUMBER



copyright © 2011 by benchware.ch

slide 18

Benchmark Results



CPU and Server: PL/SQL arithmetic mix data type NUMBER

Run	Tst	Code	#N	#S	#T	[rps]	[ops]	CPU		Read	Write	Total	Read	Write	Total	Redo	Time
								[%]	[iops]	[iops]	[iops]	[mbps]	[mbps]	[mbps]	[s]		
2	1	C231	1	1	1	0.00E+00	5.12E+05	4	8	28	36	0	0	0	0	0	39
2	C231	2	2	1	0.00E+00	1.00E+06	7	8	25	33	0	0	0	0	0	0	40
3	C231	2	4	1	0.00E+00	1.90E+06	12	10	25	35	0	0	0	0	0	0	42
4	C231	2	8	1	0.00E+00	4.00E+06	24	8	25	33	0	0	0	0	0	0	40
5	C231	2	16	1	0.00E+00	8.00E+06	48	8	26	34	0	0	0	0	0	0	40
6	C231	2	32	1	0.00E+00	1.04E+07	94	7	19	25	0	0	0	0	0	0	61
7	C231	2	64	1	0.00E+00	1.08E+07	97	7	15	22	0	0	0	0	0	0	118

Run	Tst	Code	#N	#S	#T	[rps]	[ops]	CPU		Read	Write	Total	Read	Write	Total	Redo	Time
								[%]	[iops]	[iops]	[iops]	[mbps]	[mbps]	[mbps]	[s]		
8	1	C231	1	1	1	0.00E+00	2.08E+05	9	2	9	10	0	0	0	0	0	120
2	C231	1	2	1	0.00E+00	4.17E+05	17	2	4	6	0	0	0	0	0	0	120
3	C231	1	4	1	0.00E+00	8.25E+05	33	1	4	5	0	0	0	0	0	0	121
4	C231	1	8	1	0.00E+00	1.47E+06	64	1	4	5	0	0	0	0	0	0	136
5	C231	1	16	1	0.00E+00	1.94E+06	89	1	3	4	0	0	0	0	0	0	206
6	C231	1	32	1	0.00E+00	2.11E+06	98	1	2	3	0	0	0	0	0	0	379

copyright © 2011 by benchware.ch

slide 19

Benchmark Results



CPU and Server: PL/SQL arithmetic mix data type NUMBER

	Metric	Exadata X2-2 2 DB-Server	IBM Server P-Series	Advantage Exadata	
CPU	Type	X5760	Power 7		
	Frequency [GHz]	2.93	3.5		
	#Cores per socket	6	6		
	Multithreading per Core	2	4		
Server	#Sockets per Server	2	2		
	#Cores per Server	12	12		
	Threads per Server (CPU_COUNT)	24	48		
	Oracle core based license	6	12	x 2.0	
Performance based on PL/SQL	Speed	[kops]	607	339	x 1.8
	Throughput	[kops]	9600	4812	x 2.0
	Throughput/thread	[kops]	400	100	x 4.0
	Throughput/core	[kops]	800	401	x 2.0


copyright © 2011 by benchware.ch

slide 20

Benchmark Results



Storage: random read



copyright © 2011 by benchmark.ch

slide 21

Benchmark Results



Storage Performance (1)

	Metric	Exadata V2 3 Storage Cells	Conventional Platform	Advantage Exadata
Storage Configuration	Connectivity	240 Gbit/s IB	16 Gbit/s FC	
	Frontend Cache	3 x 512 MByte	128 GByte DS 8700 12 GByte read SVC 12 GByte write SVC	x 0.01
	#HDD Disks	3 x 12 = 36	160	x 0.23
	Flash capacity	3 x 384 GByte	-	
	RAID-Level	RAID-0/1 SW mirroring	5 (7+1) HW RAID	
	Volume Manager / File system	Oracle ASM	IBM JFS2	

copyright © 2011 by benchmark.ch

slide 22

Benchmark Results



Storage Performance (2)

	Metric	Exadata V2 QR 3 Storage Cells	Conventional Platform	Advantage Exadata
Storage Performance	Sequential I/O			
	• Read Disk [mbps]	4'500	524	x 8.6
	• Read Flash [mbps]	12'000	-	x 22.9
	• Write Disk [mbps]	3'800	323	x 11.8
Random I/O				
	• Read [iops]	288k	38k	x 7.6
	• Write [iops]	-	-	-


copyright © 2011 by benchware.ch

slide 23

Benchmark Results



Database: sequential data read



copyright © 2011 by benchware.ch

slide 24

Benchmark Results



Database Performance (1)

Oracle 11.-2 Enterprise Edition

	Metric	Exadata V2 QR 3 Storage Cells	Conventional Platform	Advantage Exadata
Data Load Conventional				
• DOP = 1 (speed)	[rps]	7'400	4'170	x 1.8
• DOP = 4 (throughput)	[rps]	22'400	8'930	x 2.5
Data Load Bulk Uncompressed				
• DOP = 1 (speed)	[rps]	360k	276k	x 1.3
• DOP = 8 (throughput)	[rps]	1'384k	711k	x 1.9
Data Load Bulk Compressed OLTP				
• DOP = 1 (speed)	[rps]	136k	133k	x 1.0
• DOP = 16 (throughput)	[rps]	2'572k	1'130k	x 2.3

Benchmark Results



Database Performance (2)

Oracle 11.-2 Enterprise Edition

	Metric	Exadata V2 QR 3 Storage Cells	Conventional Platform	Advantage Exadata
Data Scan Sequential				
• Buffer Cache (SGA)	[mbps]	-	-	-
• Disk	[mbps]	4'400	560	x 7.9
• Cell Flash Cache	[mbps]	12'000	(560)	X 21.4
Data Select Random				
• Cached	[rps]	365k	161k	x 2.3
• Non-Cached	[rps]	299k	42k	x 7.1
	[iops]	288k	41k	x 7.0
Data Update Random				
• Cached	[rps]	45k	40k	x 1.1
• Non-Cached	[rps]	8k	6k	x 1.3
	[iops]	8k	7k	X 1.1

Contents



- 1 Issues with proof-of-concepts
- 2 Complexity of database platforms
- 3 Useless TPC benchmarks
- 4 The Benchware approach: Key Performance Metrics
- 5 Benchmark Example: conventional platform versus Exadata
- 6 Conclusions

Conclusions

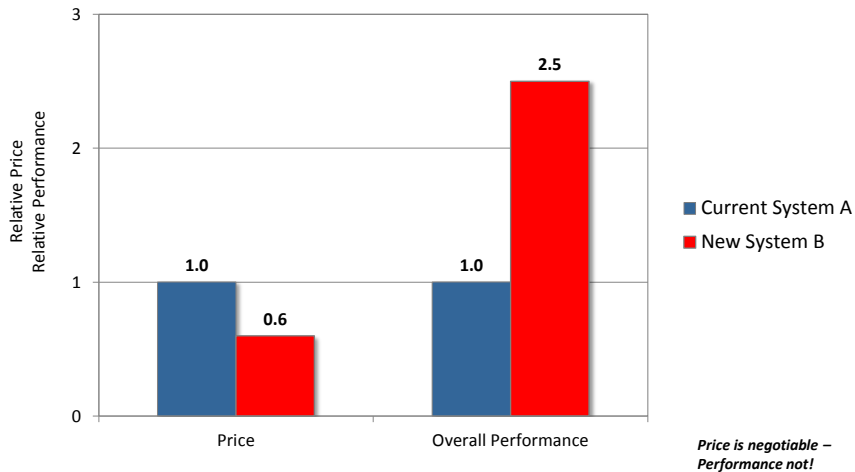


- Few key performance metrics describe overall database platform performance
- Simple benchmark determines key performance metrics to compare price performance ratio
 - Architectures
 - Components
 - Configurations
- Also useful for new engineered platforms
 - Calibrate and certify Oracle platform performance
 - Identify performance bottlenecks of Oracle platform

Conclusions



Price-/Performance ratio based on customers priorities



copyright © 2011 by benchware.ch

slide 29

Key value proposition



Benchware approach

- Benchmark platform to determine performance capabilities
 - useful for capacity planning and service level agreements
- Comprehensive – all basic database operations
- Simple – to install and to use
- Fast – to run, within a few hours
- Scalable – from laptop to high-end system
- Reproducible
- Understandable performance metrics
- Uses Oracle features – like RAC, DataGuard, Flash Cache, Exadata, ...

copyright © 2011 by benchware.ch

slide 30

About the author



Manfred Drozd
Dipl.-Inform.

Manfred Drozd studied Computer Science at the University of Paderborn (Germany). He observed the relational database technology from the beginning when he started his career in 1980 as a programmer developing a relational database system. A life science company in Basle hired him in 1984 to implement Oracle Version 3.1 at the R&D data center. During that time he also lectured courses in Computer Architecture and Database Systems at the HTL in Berne and Basle. Between 1986 and 1990 he managed several database development teams. From 1990 to 2001 Manfred Drozd was employee of Oracle Corp. Switzerland, lastly founding and heading the consulting practice *Server Technology & Performance Architecture*. Currently he is working as an independent Consultant designing, implementing, benchmarking and optimizing Oracle database platforms.

Since 1995 Manfred Drozd focuses on Oracle performance and architecture. On behalf of customers he periodically runs performance tests in the benchmark centers of the hardware vendors. He also lectures training courses and public seminars about scalable Oracle Systems and Oracle performance tuning. He is frequently speaker at SOUG (Swiss Oracle User Group) and DOAG (Deutsche Oracle Anwendergruppe) events. Manfred Drozd and his team have developed benchmark tools over the last years to identify key performance metrics of Oracle platforms. Benchmarking helps to understand platform performance based on factual knowledge.

Manfred Drozd is advocator of a holistic *Performance by Design* approach: Oracle Database platforms are build from bottom up with a complete calibration of all technology layers focusing on the performance and availability requirements of applications. He used this approach very successfully for the architecture of large OLTP and Data Warehouse Systems in the telecommunication and financial industry.

BENCHWARE

swiss precision in performance measurement

www.benchware.ch

info@benchware.ch