BENCHWARE

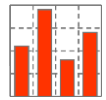

Oracle auf HP/Violin – Wirklich ein Exadata Killer?

DOAG Jahreskonferenz

November 2012

Contents

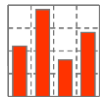

Benchware Ltd

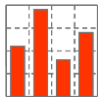

Services and Products

Strong foundation in core technologies such as
Oracle database systems, server and storage systems

- System Architecture, Component Evaluation, Reviews
- Performance Analysis & Optimization
- Benchmarking
- Database engineering

Benchware Ltd

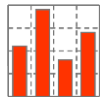

Value proposition

- Vendor-independent company
  - Benchware is completely committed to customers' interests
- Holistic approach in designing, tuning and benchmarking Oracle systems
- Long experience track record
  - Responsible for system architecture of largest DWH and OLTP systems, mainly telecom and finance industry
  - Oracle since 1984 (Oracle Version 3)
  - Performance tuning and benchmarking since 1993 (Oracle Version 7)

Performance of complex Oracle platforms is not predictable

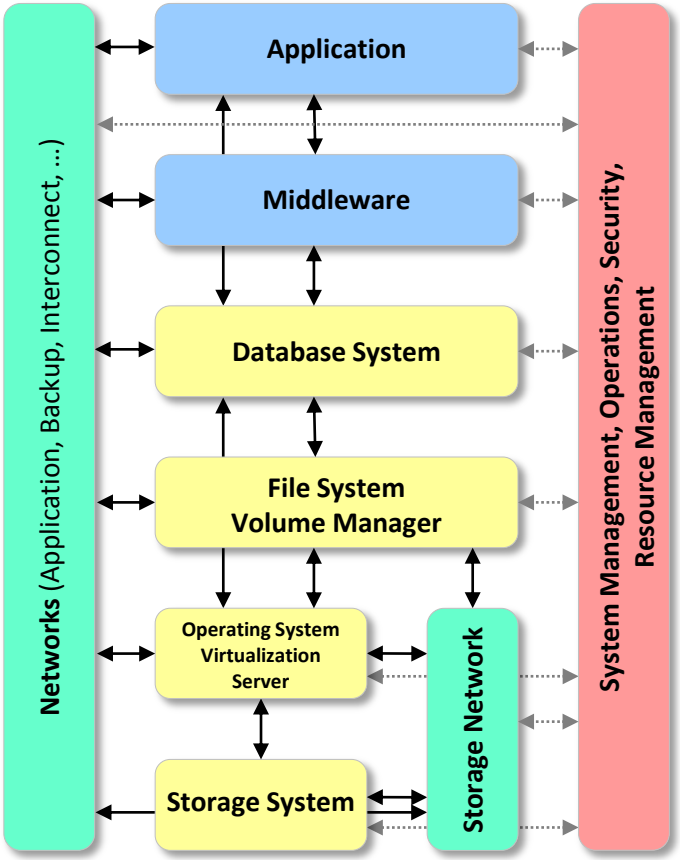

System architects have a wide choice of components, technologies and configurations

Networks (IP-based)

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

Oracle Database

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

Storage Network (FC-, IB- or IP-based)

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 configuration parameters (#LUNS, 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 one 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.

Benchware Ltd

Complex architecture of Oracle platforms requires benchmarking

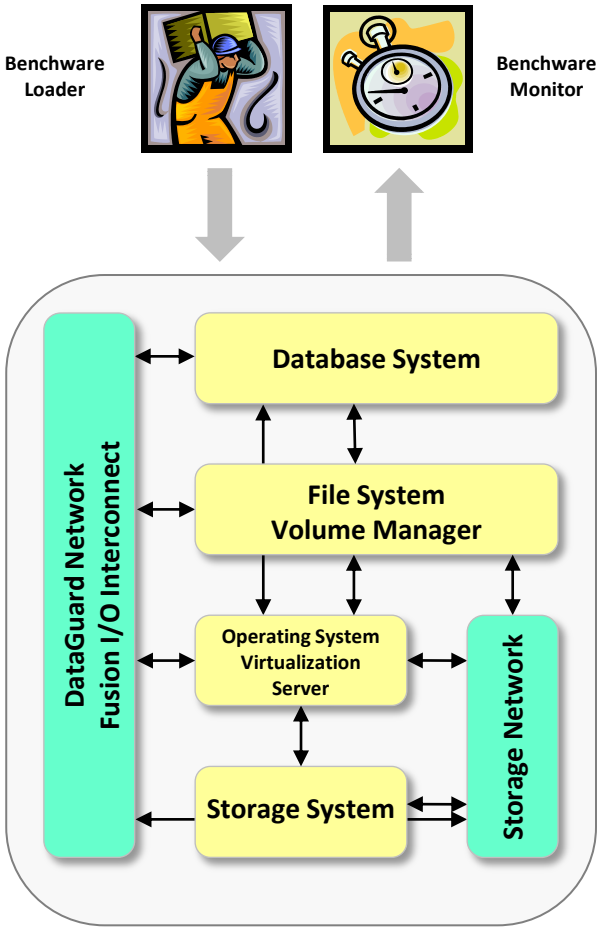

Object of measurement: Oracle platform

- Benchware Performance Suite
  - Benchware Monitor
  - Benchware Loader
- Performance measurement at the interface between application and Oracle database platform
- Key Performance Metrics can be used for SLA between IT operation and business
- Benchware uses Oracle database software to generate all kinds of loads for cpu, server, storage and database

Benchware Ltd

Key performance metrics of Oracle platforms should be as easy as key performance metrics used in the automobile industry

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

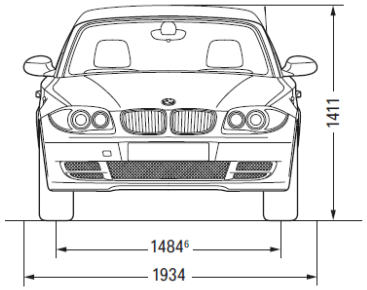

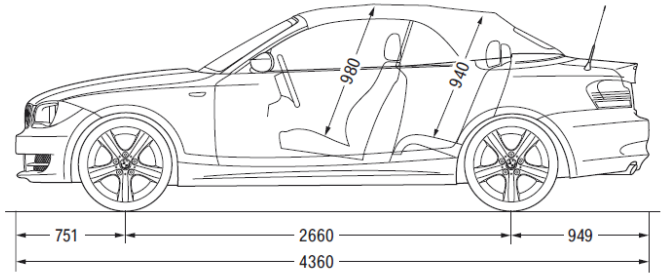

Benchware Ltd

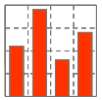

Library of Oracle benchmark tests - implemented in PL/SQL, Java and SQL

CPU Performance CPU-bound PL/SQL operations	OLTP systems	DWH systems	Proof of Efficiency	Key Performance Metrics	Unit
• pl/sql basic operations arithmetic mix, string mix	★★	★★	multithreading virtualization	speed throughput	[s] [ops]
• pl/sql algorithms fibonacci, prime numbers	★★★	★★			

Server Performance Server-bound SQL operations completely in SGA	OLTP systems	DWH systems	Proof of Efficiency	Key Performance Metrics	Unit
• in-memory SQL full scan, index access	★★★	★★	scalability cc-numa virtualization	speed throughput	[µs] [s] [bgps] [tps] [rps]
• pl/sql algorithms quicksort	★★	★			

[s]	seconds	[bgps]	buffer gets per second	[MBps]	mega bytes per second	★	less important
[ms]	milli seconds (10^{-3})	[rps]	rows per second	[GBps]	giga bytes per second	★★	important
[µs]	micro seconds (10^{-6})	[tps]	transactions per second	[iops]	i/o operations per second	★★★	very important
[ns]	nano seconds (10^{-9})	[ops]	operations per second	[qpm]	queries per minute		

Benchware Ltd

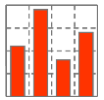

Library of Oracle benchmark tests - implemented in PL/SQL, Java and SQL

Storage Performance I/O-bound Oracle operations	OLTP systems	DWH systems	Proof of Efficiency	Key Performance Metrics	Unit
• sequential I/O 1 MByte, read and write	★★	★★★	data integrity tiering, pooling striping virtualization replication	throughput service time	[ms] [MBps] [GBps] [iops]
• random I/O 8 kByte, read and write	★★★	★			

Database Performance Mixed resource usage: CPU, memory, storage	OLTP systems	DWH systems	Proof of Efficiency	Key Performance Metrics	Unit
• data load uncompressed, compressed	★★	★★★	scalability	speed throughput service time	[ms] [s] [rps] [tps] [qpm] [s]
• data scan	★	★★★			
• data aggregation & reports	★★	★★★			
• OLTP transactions insert, select, update	★★★	★			

[s]	seconds	[bps]	buffers per second	[MBps]	mega bytes per second	★	less important
[ms]	milli seconds (10^{-3})	[rps]	rows per second	[GBps]	giga bytes per second	★★	important
[µs]	micro seconds (10^{-6})	[tps]	transactions per second	[iops]	i/o operations per second	★★★	very important
[ns]	nano seconds (10^{-9})	[ops]	operations per second	[qpm]	queries per minute		

Benchware Ltd

Public benchmark results: www.benchware.ch/benchmarks

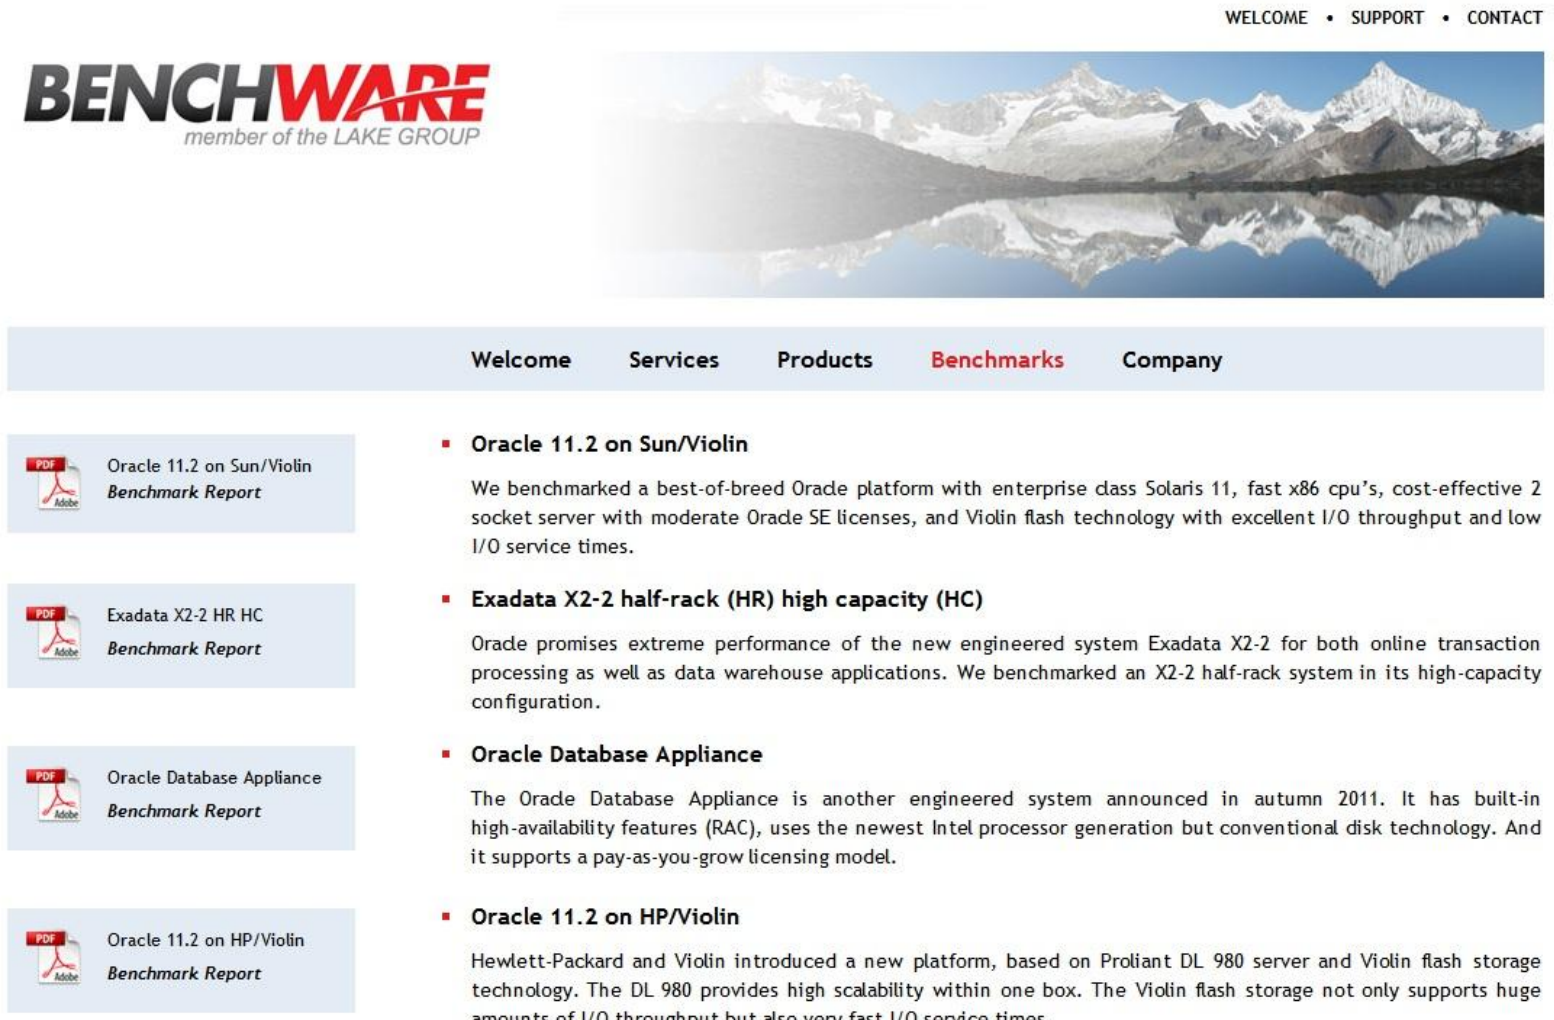

Contents

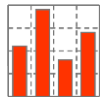

Architecture Overview

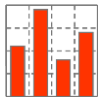

Flash Technology will change computer system architecture dramatically

- Hardware vendors offer different architectures with flash technology for Oracle platforms
- The architectures provided by hardware vendors differ very much
  - Volatility
  - Capacity
  - Throughput
  - Latency
  - Cost (per GByte, per IOPS)
  - Manageability
  - Shareability

Conventional System Architecture

Architecture without any Flash Technology

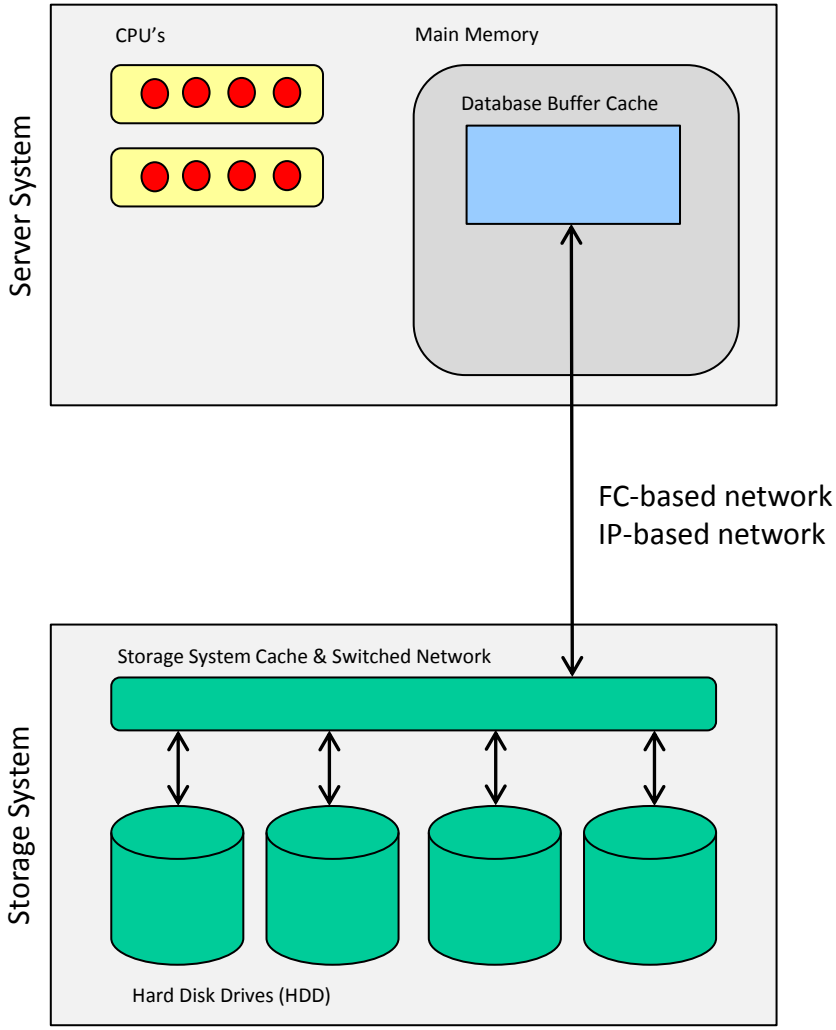

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s

- Storage system cache 1×10^{-3} s
- Storage system disk 1×10^{-2} s

Storage System with solid state disks

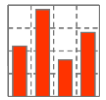

Automatic tiering, e.g. EMC, HDS or IBM high-end storage systems

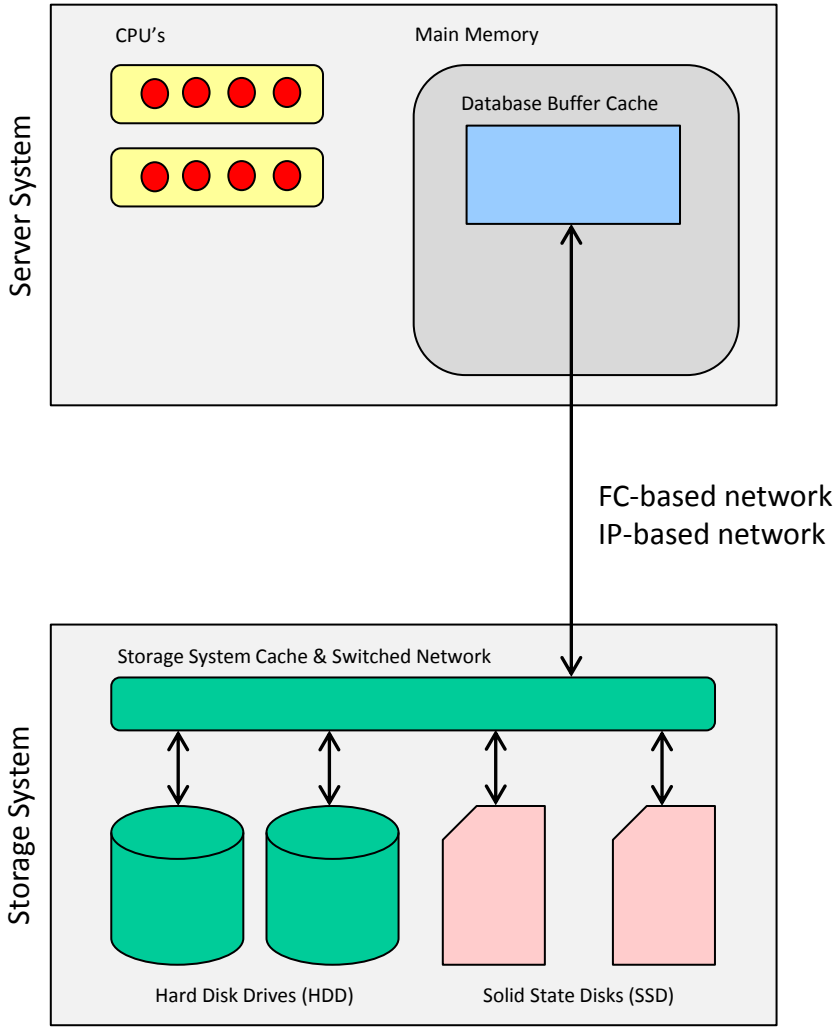

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s
- Solid State Disk 500×10^{-6} s
- Storage system cache 1×10^{-3} s
- Storage system disk 1×10^{-2} s

Storage System based on Flash Technology

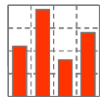

Storage Systems completely build on flash technology, e.g. Violin

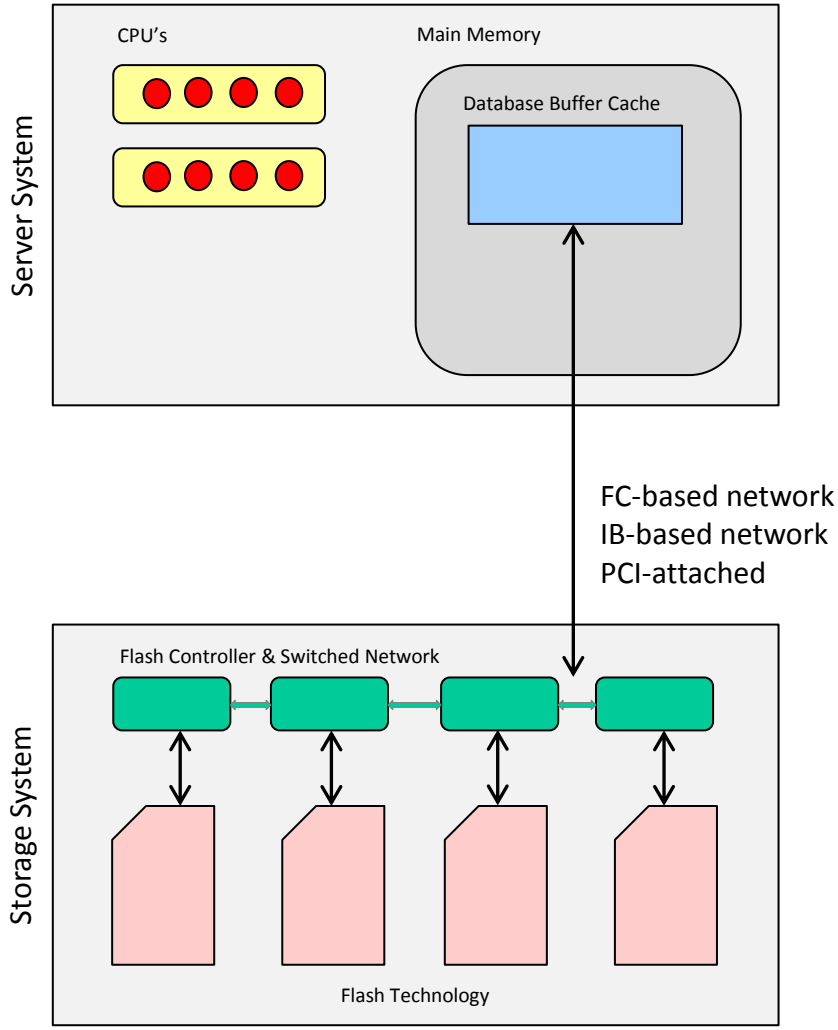

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s
- Flash Disk 250×10^{-6} s

Oracle Flash Cache Technology

Specific solution only for Solaris and OEL and Sun F20 Flash

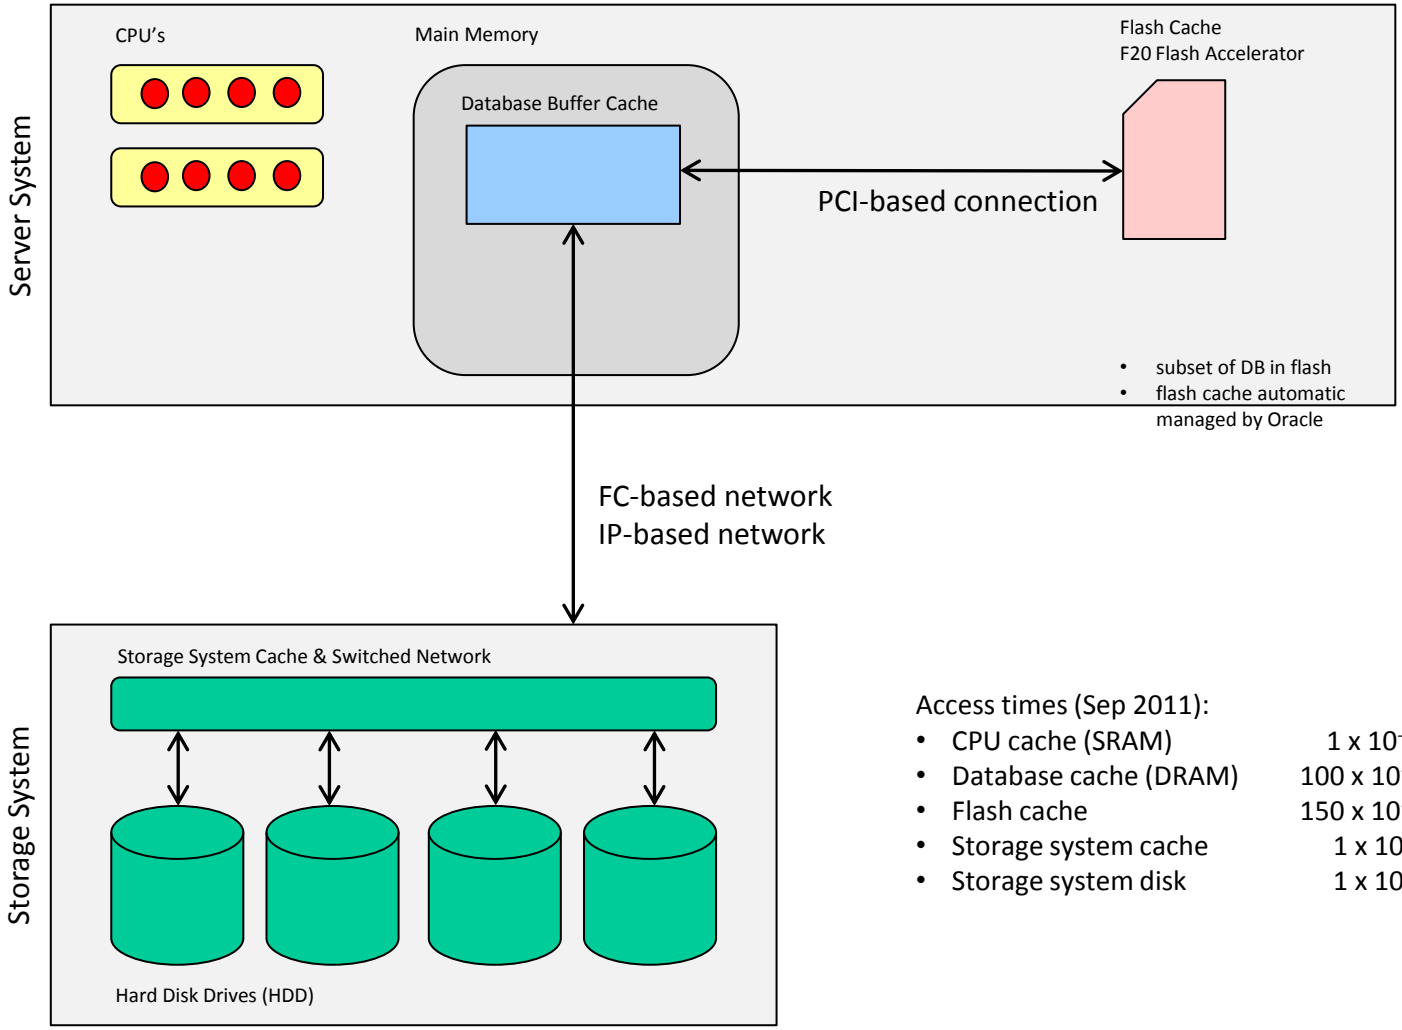

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s
- Flash cache 150×10^{-6} s
- Storage system cache 1×10^{-3} s
- Storage system disk 1×10^{-2} s

Mirrored Database in Server Fusion IO Cards

Mirrored ASM failure group, e.g. Hitachi Converged Platform

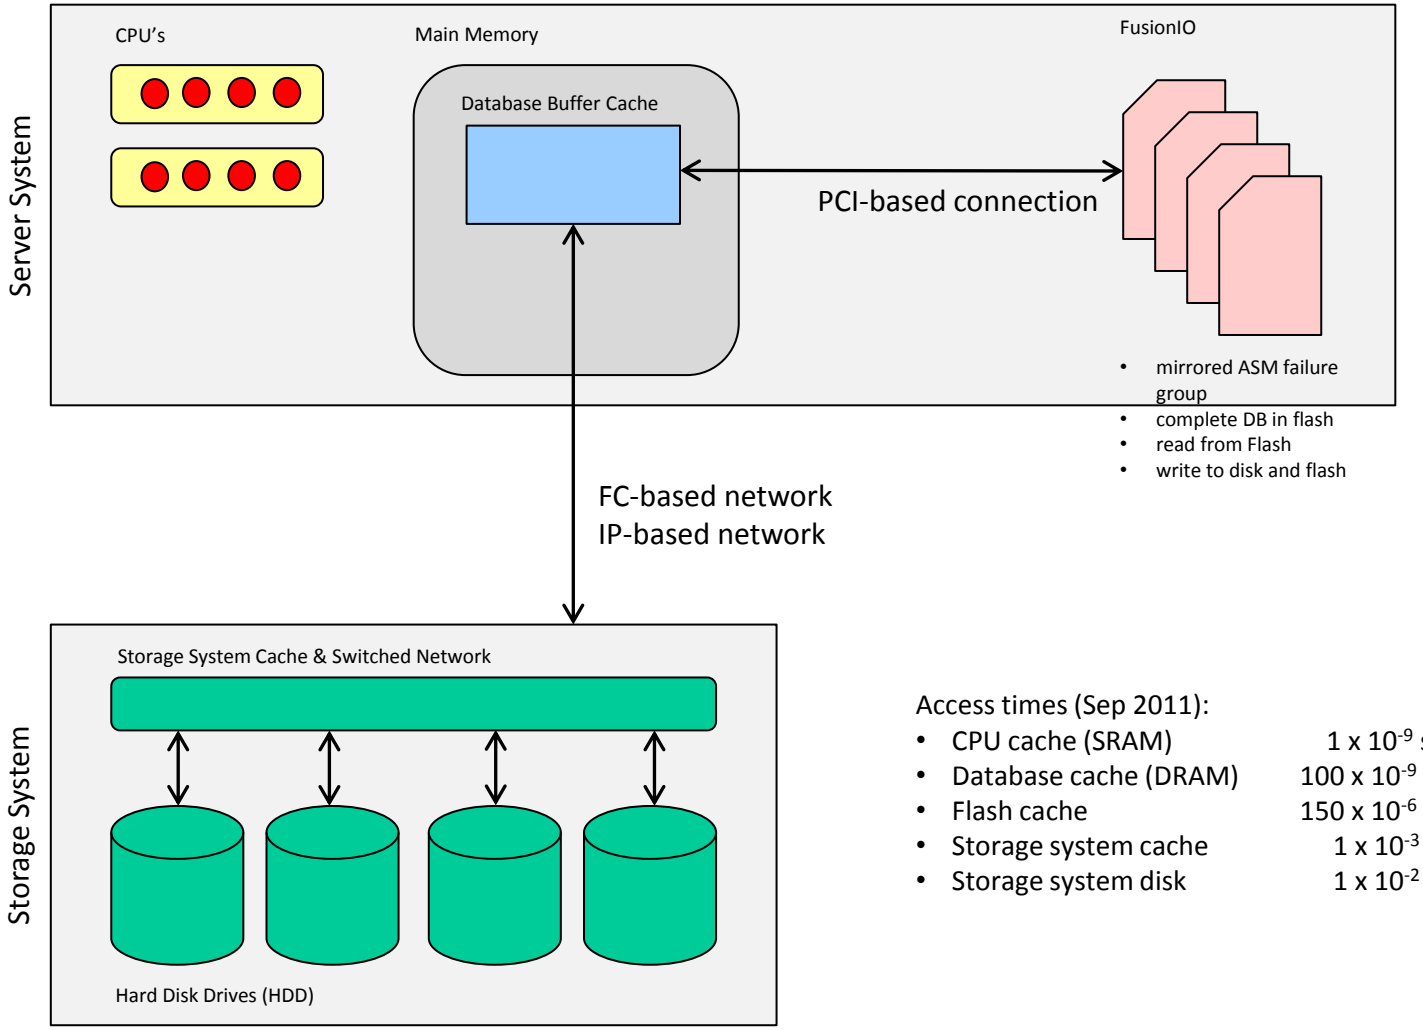

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s
- Flash cache 150×10^{-6} s
- Storage system cache 1×10^{-3} s
- Storage system disk 1×10^{-2} s

Intelligent Storage Cell

Offloaded database functions, e.g. Oracle Exadata

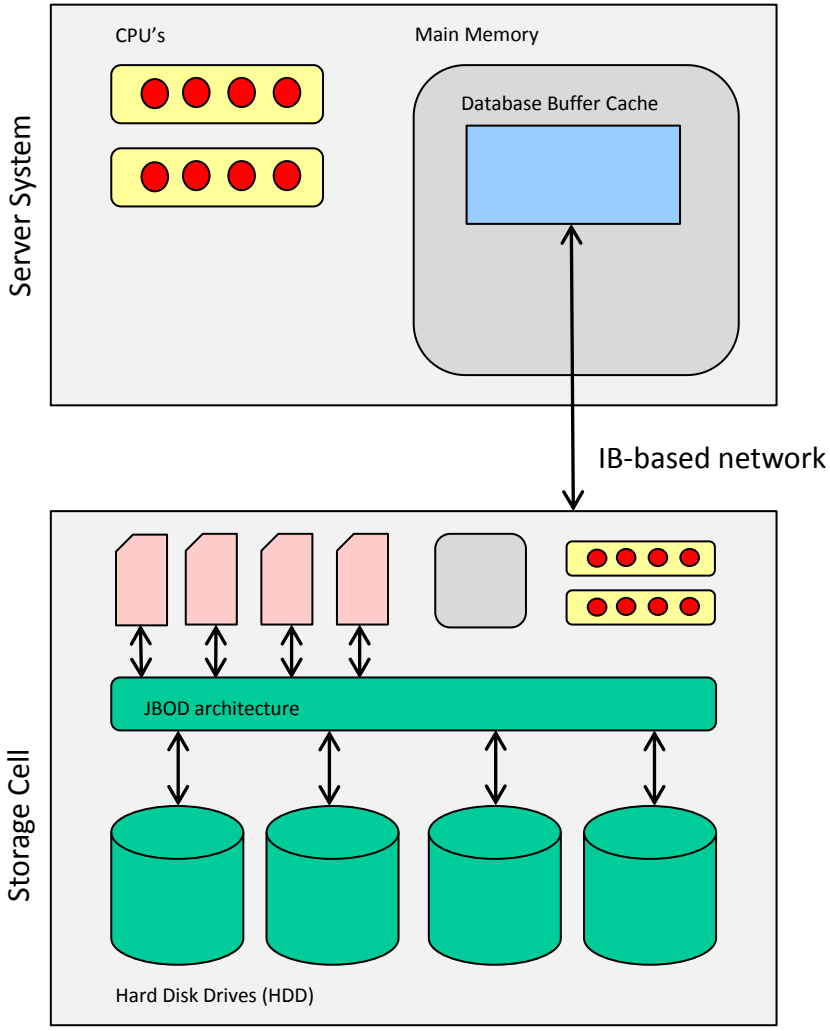

Access times (Sep 2011):

- CPU cache (SRAM) 1×10^{-9} s
- Database cache (DRAM) 100×10^{-9} s
- Flash cache 250×10^{-6} s
- Storage system cache 1×10^{-3} s
- Storage system disk 1×10^{-2} s

Contents

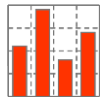

Qualitative Comparison

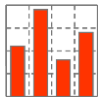

Engineered Systems

- Exadata
  - Completely pre-installed
  - No engineering
  - No tuning

- HP/Violin
  - Reference Architecture
  - High flexibility
  - Limited partnership of HP and Violin (3PAR)
  - Difficult to get official documentation www.flashdba.com

Qualitative Comparison

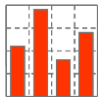

Vendor lock-in

- Exadata
  - Industry standard components
  - Hybrid-Columnar Compression

- HP/Violin
  - Industry standard components
  - Standard Oracle compression technology (not column based)

Qualitative Comparison

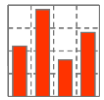

CPU- and Server

- Exadata
  - x86 offers best price-/performance ratio
  - X3-2 fast 2 socket server
  - X3-8 scalable 8 socket server
  - Larger increments
  - New 1/8 rack addresses this problem

- HP/Violin
  - x86 offers best price-/performance ratio
  - High flexibility with HP Proliant family
  - Smaller increments

Qualitative Comparison

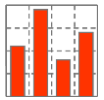

Scalability: shared disk versus shared memory

- Exadata
  - Application parallelism needed
  - Real Application Cluster needs specific programming for scalability

- HP/Violin
  - Application parallelism needed
  - Nearly no changes to application programming

Qualitative Comparison

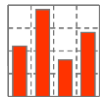

Flash Technology

- Exadata
  - Hybrid architecture
  - HDD for capacity
  - Flash for performance
  - Automatic tiering
  - Manual configuration possible
  - Writable flash since 11/2012
  - Offloading function

- HP/Violin
  - Flash only
  - Highest throughput and shortest service times for ALL database objects

Qualitative Comparison

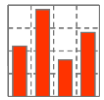

High Availability and Desaster Recovery

- Exadata
  - Built-in Real Application Cluster
  - Engineering necessary: Data Guard or Golden Gate

- HP/Violin
  - Engineering necessary
  - Engineering necessary: Host-based mirroring (ASM) or database-based mirroring (DataGuard or GoldenGate)

Qualitative Comparison

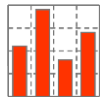

Maintenance and Support

- Exadata
  - Less complex
  - One Patch for complete technology stack

- HP/Violin
  - More complex
  - Patches for each product from each vendor

Contents

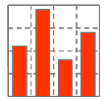

Quantitative Comparison

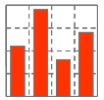

Selected benchmark numbers will be presented during the presentation on 22nd November 2 p.m.

Contents

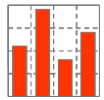

Conclusion

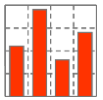

- Both solutions provide individual strengths and weaknesses
- HP/Violin offers more flexibility for the price of more engineering and error-proneness
- Exadata provides standardization and built-in performance by offloading functions, which can not be exceeded by competition
- To answer the question: there is currently no Exadata Killer around

swiss precision in performance measurement

www.benchware.ch

info@benchware.ch

Benchware Ltd

For more information please contact

Manfred Drozd
Dipl.-Inform.
Co-founder & Managing Partner

manfred.drozd@benchware.ch
+41 79 334 68 68

Benchware AG
Seestrasse 18
CH-8800 Thalwil (Switzerland)

info@benchware.ch
+41 44 722 16 16
www.benchware.ch

Manfred Drozd studied Computer Science at the University of Paderborn (Germany). He observed the relational database technology from its beginnings 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 an employee of Oracle Corp. Switzerland, ultimately 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 1993 Manfred Drozd has been focusing on Oracle performance and architecture. On behalf of customers he periodically runs performance tests in the benchmark centers of the hardware vendors. He also holds training courses and public seminars about scalable Oracle systems and Oracle performance tuning. He is a frequent speaker at SOUG (Swiss Oracle User Group) and DOAG (Deutsche Oracle Anwendergruppe) events. Over the past 12 years Manfred Drozd and his team have developed benchmark tools to identify Oracle platform key performance metrics. Benchmarking helps to understand platform performance based on factual knowledge.

Manfred Drozd is an advocator of an holistic *Performance by Design* approach: Oracle Database platforms are built from the 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.