

HP Universal Solutions for Oracle



Dr. Draško Tomić
HP Chief Technologist

16/10/2013. Rovinj, Croatia



Agenda



Why HP solutions for Oracle?

HP Universal Solutions for Oracle

About the Performance & Endurance

Some benchmarking results

Q & A



Why HP Solutions for Oracle?

CO18
HROUG




© Copyright 2012 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

The portfolio of Oracle EXA machines

Big data appliance



Exalytics High-speed analytics



Exadata

Oracle database for:

- Online transaction processing (OLTP)
- Business intelligence (BI)
- Database (DB) consolidation

Exalogic

Fusion middleware and applications

Main Exadata challenges

- Performance
- Scalability
- Upgradeability
- Open Source issues
- Price




HP solutions for Big Data

HP AppSystem for Apache



HP AppSystems for SAP HANA

Delivering end-to-end insights



HP AppSystem for Microsoft SQL Server 2012 for Parallel Data Warehouse



HP Universal Database Solution for MC x86



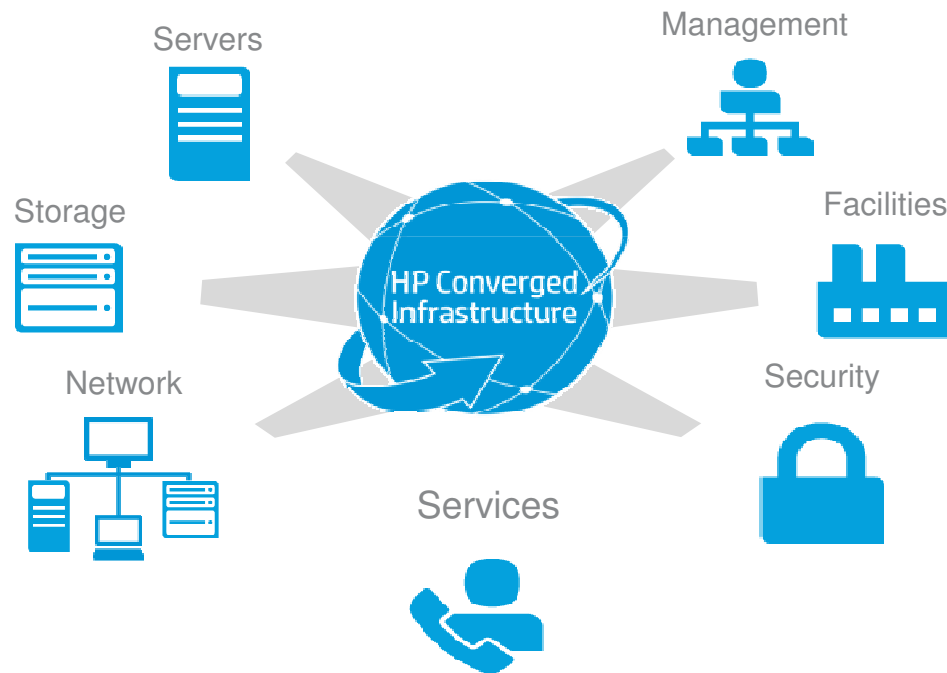
HP Moonshot – solving the Biggest Big Data Challenges

The world's first software defined server



HP Converged Infrastructure – the foundation for HP Big Data Solutions

CO18
HROUG



Customers protect their investment



HP solutions are built on HP's Converged Infrastructure open standards-based systems


- Supports heterogeneous environments
- Not locked into a proprietary stack—choose best-of-breed
- HP innovations; Gen8 servers, Project Moonshot, Project Odyssey, Project Voyager, Virtual Application Networks—and the list goes on
- Tier 1 storage, thin provisioning, optimized tiered storage, multi-tenancy, and full integration within existing data center



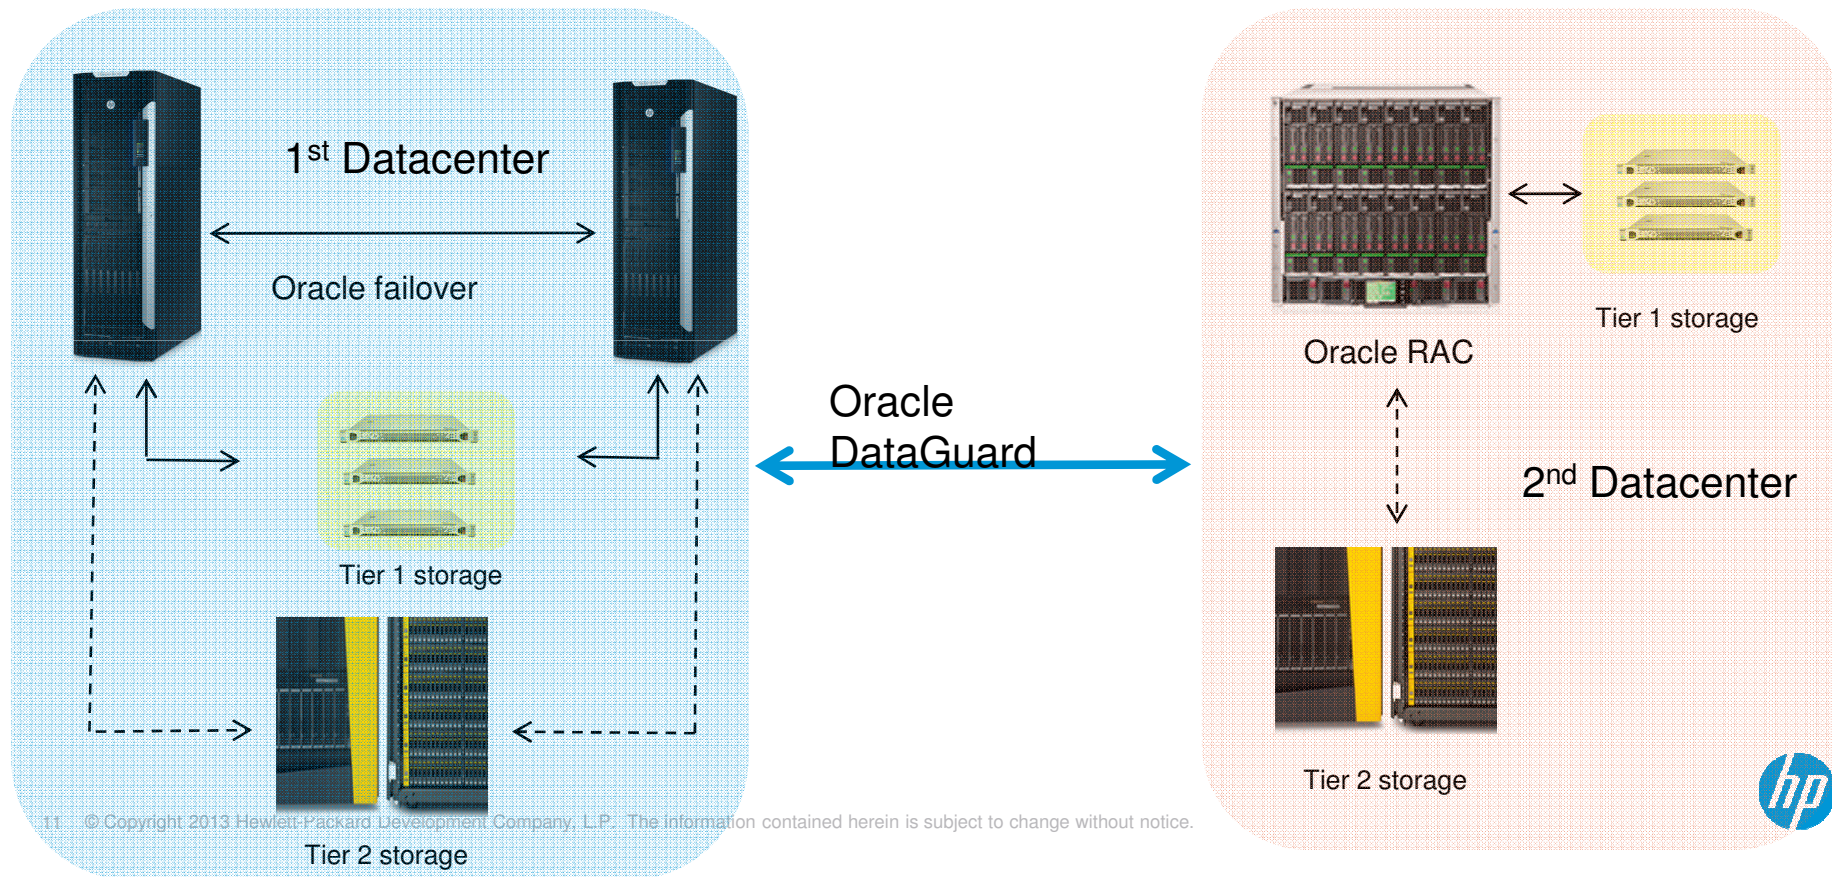
HP Universal Solutions for Oracle



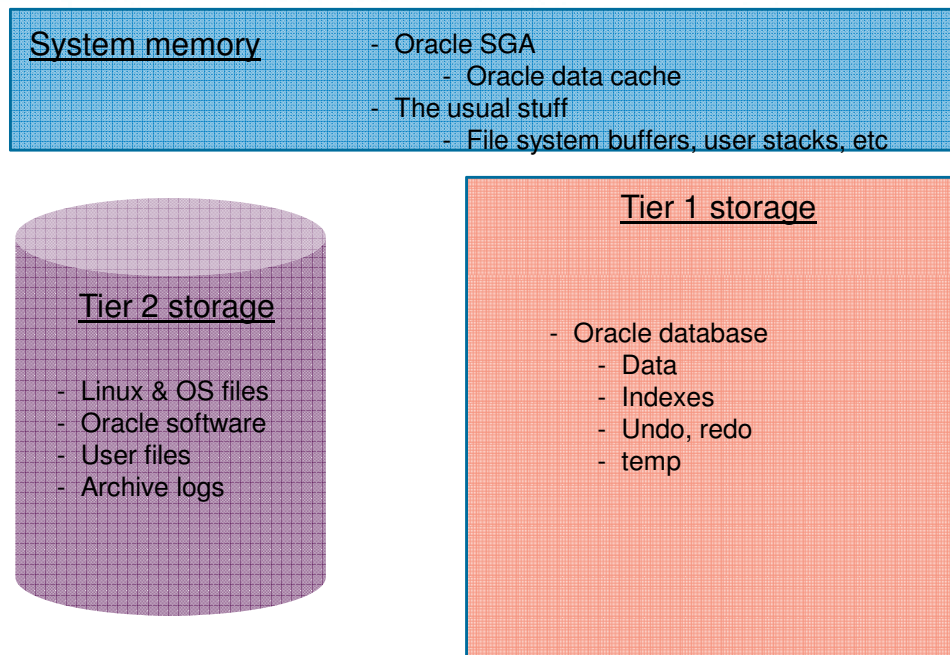
Simplicity is the ultimate sophistication
—Leonardo da Vinci



The Framework



Tiered storage



Tier 1 storage



Hardware: HP Proliant DL380/560/580 servers

Interface: FC or Infiniband

Accelerators: FUSION-io

Acceleration: shared or virtualized

Software: ION Data Accelerator, ioTurbine,
ioVDI

High-availability: Oracle ASM

Disaster recovery: Oracle DataGuard




HP Proliant DL380 Gen8 servers



Tier 1 examples

- DL380 Gen8 server
 - 2 x Ivybridge CPUs
 - 2 x dual-port 16Gb FC
 - 4 x 1.2TB flash IO cards.

- DL580 Gen7 server
 - 4 x Westmere CPUs
 - 2 x dual-port 16Gb FC
 - 8 x 2.4TB flash IO cards



Tier 2 storage



Storage: HP 3PAR 7x00 storage system

Interface: FC

Drives: SAS/SATA

Software: 3PAR Adaptive Optimization



HP 3PAR Storage System




More on integration with Oracle

High Availability? No

problem.

- Oracle Data Guard
 - Already paid for with site license
 - RAC not required



- Data Guard configuration consisting of a primary database and a physical standby database. From the primary database, redo is being transmitted and applied to the standby database. Log apply services apply the redo out of the standby redo log files to the standby database.

A Perfect fit with Oracle ASM

- NAND Flash drives appear as raw devices
- Oracle ASM can consume the NAND raw devices to create Disk Groups and Failure Groups
- Use of ASM provide great use cases for I/O balancing
- ASM provides great levels of redundancy
- Points of failure are minimized

Single instance Oracle implementations

Can use with ASM or without ASM

Pure Flash implementation


- Place entire database on Flash

Hybrid Flash implementation


- Place redo logs on Flash
- Place specific tables/indexes on Flash
- Use Partitioning to locate a portion of data of a table to Flash
- Oracle Smart Flash Cache

Accelerator types



PCI single memory module



PCI dual memory module




Blade mezzanine cards



Accelerator MLC portfolio for HP Proliant servers

18
HROUG

<p>365GB mezzanine</p> <p>Read IOPS: 415.000 Write IOPS: 530.000</p>	<p>HP 365GB PCIe</p> <p>Read IOPS: 415.000 Write IOPS: 535.000</p>
<p>785GB mezzanine</p> <p>Read IOPS: 443.000 Write IOPS: 530.000</p>	<p>HP 785GB PCIe</p> <p>Read IOPS: 443.000 Write IOPS: 535.000</p>
<p>1.2 TB mezzanine</p> <p>Read IOPS: 443.000 Write IOPS: 530.000</p>	<p>HP 1.2TB PCIe</p> <p>Read IOPS: 445.000 Write IOPS: 535.000</p>
	<p>HP 2.4TB PCIe</p> <p>Read IOPS: 892.000 Write IOPS: 935.000</p>
	<p>HP 3TB PCIe</p> <p>Read IOPS: 445.000 Write IOPS: 535.000</p>



Endurance



SLC (1st generation) modules are lower density (less data stored per unit area), but consume less power, can operate at higher temperatures, exhibit faster write/erase speeds, and endure a significantly greater number of write/erase cycles than MLC modules.

MLC (2nd generation) modules allow for the highest storage density at the lowest cost. MLC memory is known to be subject to a higher rate of raw bit errors, but the 39-bit error correction in IO Accelerators guards against such errors.

Usage of sophisticated algorithms to ensure modules are uniformly written with data.



Maximum usage – amount data written before burning-out



SLC

320GB: 4PB

640GB: 8PB

MLC

365GB:

4PB

785GB:

11PB

1205GB:

17PB

2410GB:

34PB

Example: 2410GB module written with 18630GB daily will burn-out after 5 year period

3000GB:

37PB



Wrap-up



Typical configuration is x86 DL380 or 560 servers with Fusion-io cards.

Fusion-io has software components (Directcache, or IO Turbine) that use the Fusion-io card as a read cache in front of the general storage array.

This gives us a very fast IO. If the application is small enough, we can use 100% Fusion-io (10TB per server is ok).

Also it is recommended to use a lot of memory, even max. configuration (0,5 -1 TB) is cheap enough today.

The trick is that most Oracle applications require a lot of memory and fast IO, but not too much CPU. So we can MINIMIZE the CPU core count for the system. As Oracle is licensed per CPU, and it costs much more than the HW, we can save big amount of Oracle license this way and have much better 5 years TCO altogether.

If the customer requires bigger/more complex environment, we can also use 40Gbit Infiniband as interconnect between servers.



Software from Fusion-

io Data Accelerator

- Integrated with Oracle Enterprise Manager and Oracle VM Server
- Plugin for Oracle 11g and 12c database to monitor performance and system status.
- Certified for Oracle VM Server

ioVDI

- data fetched from server-side flash
- intelligent off-load of most reads and up to 80% writes from tier 2 storage layer
- efficient use of CPU & I/O resources for 100s of desktops hosted on a single server

ioTurbine

22 © Copyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

- option for hypervisor caching




Some Benchmarking Results



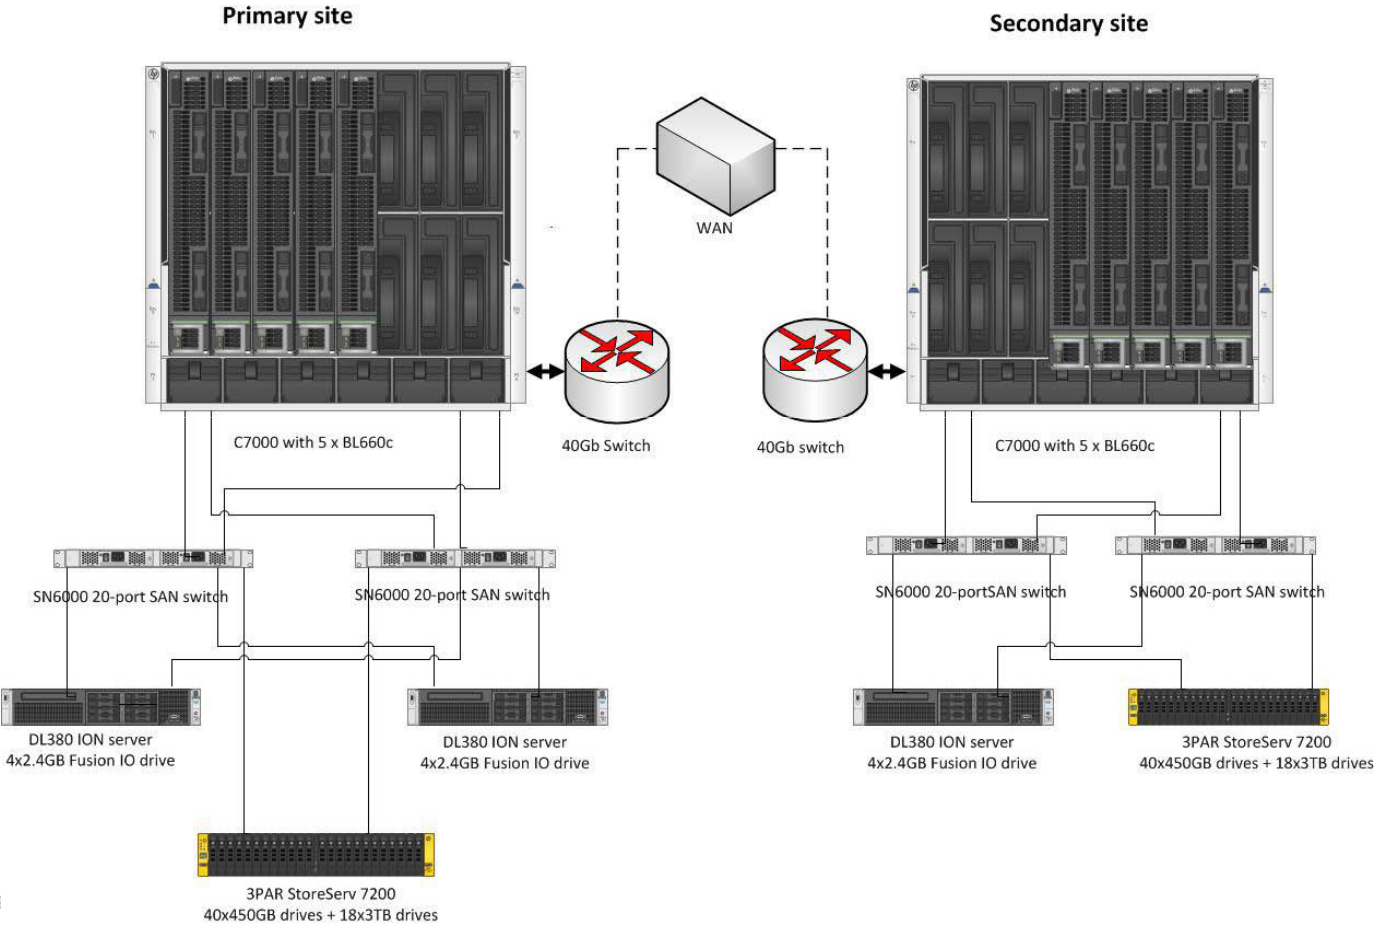
Benchmarking case #1 (as of 13th Sept 2013)



HP achieves world records for 2P blade and 2P rack-mount server virtualization performance on VMmark 2.5.1 benchmark



Benchmarking case #2 (as of 7th August 2013)



Components used

DB tier:

5 x HP Proliant BL660c Gen8 RAC cluster, with 4 x E5-4650 CPUs, 12 x 16GB RAM modules, 2 x 8Gb dual-port FC HBAs, 2 x 10Gb dual-port LAN adapters, 1GbE LOM.


Storage tier:

2 x HP Proliant DL380p Gen8 flash storage (tier 1 storage)

DL380p: 2 x E5-2690 CPUs, 12 x 8GB RAM modules, 2 x 16Gb dual-port FC HBAs, 4 x 2.4GB MLC IO Accelerators.

Software used:

RHEL5 operating system
Oracle 11.2.0.3.0 EE
ION Data Accelerator
HammerDB (aka HammerOra)



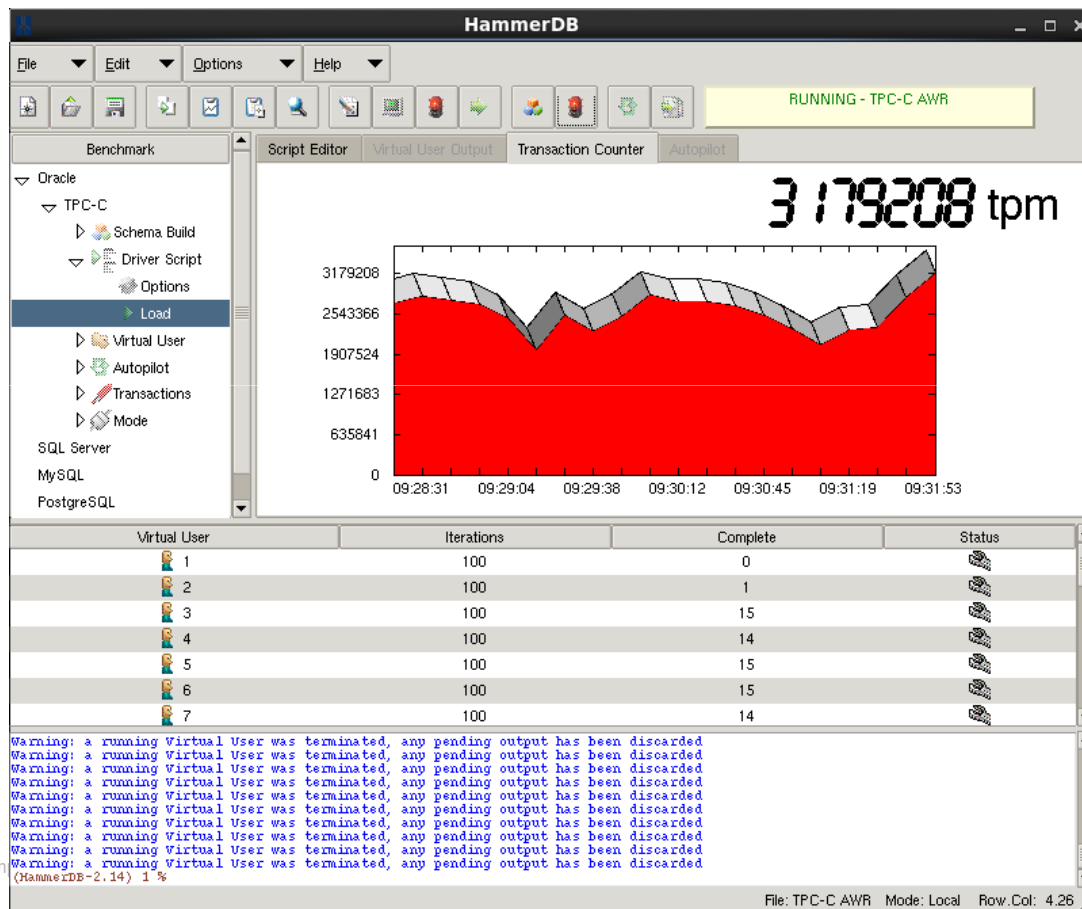
Results – transactions per minute




Up to 3.2M tpms

2.3M tpms sustained

9M tpms peak
(estimated)



Results – reads/writes



Results – IOPS

