

# Update: Testing Exadata, the new Oracle database machine

9 December 2008

Anton Topurov

Eric Grancher

Elzbieta Gajewska-Dendek



- Oracle Database Machine
- Exadata Storage Server Hardware
- Exadata Features
- Beta Testing
- Recent Tests
- Conclusions

# Oracle Database Machine




- 8 DL360 Oracle Database servers
  - 2 quad-core Intel Xeon, 32GB RAM
  - Oracle Enterprise Linux
  - Oracle RAC
- 14 Exadata Storage Cells (SAS or SATA)
  - Up to 14 TB uncompressed user data on SAS
  - Up to 46 TB uncompressed user data on SATA
- 4 InfiniBand switches
- 1 Gigabit Ethernet switch
- Keyboard, Video, Mouse (KVM) hardware
- Hardware Warranty
  - 3 YR Parts/3 YR Labor/3 YR On-site
  - 24X7, 4 Hour response time

## Exadata Storage Server



## Racked Exadata Storage Servers



- **Building block of massively parallel Exadata Storage Grid**
  - Up to 1GB/sec data bandwidth per cell
- **HP DL180 G5**
  - 2 Intel quad-core processors
  - 8GB RAM
  - Dual-port 4X DDR InfiniBand card
  - 12 SAS or SATA disks
- **Software pre-installed**
  - Oracle Exadata Storage Server Software
  - Oracle Enterprise Linux
  - HP Management Software
- **Hardware Warranty**
  - 3 YR Parts/3 YR Labor/3 YR On-site
  - 24X7, 4 Hour response

Database aware storage – does:

- Predicate filtering
- Column projection filtering
- Join processing (star-joins for DWH)
- ***Tablespace creation***
  - eliminates the I/O associated with the creation and writing of tablespace blocks
- I/O resource management – inter and intra database

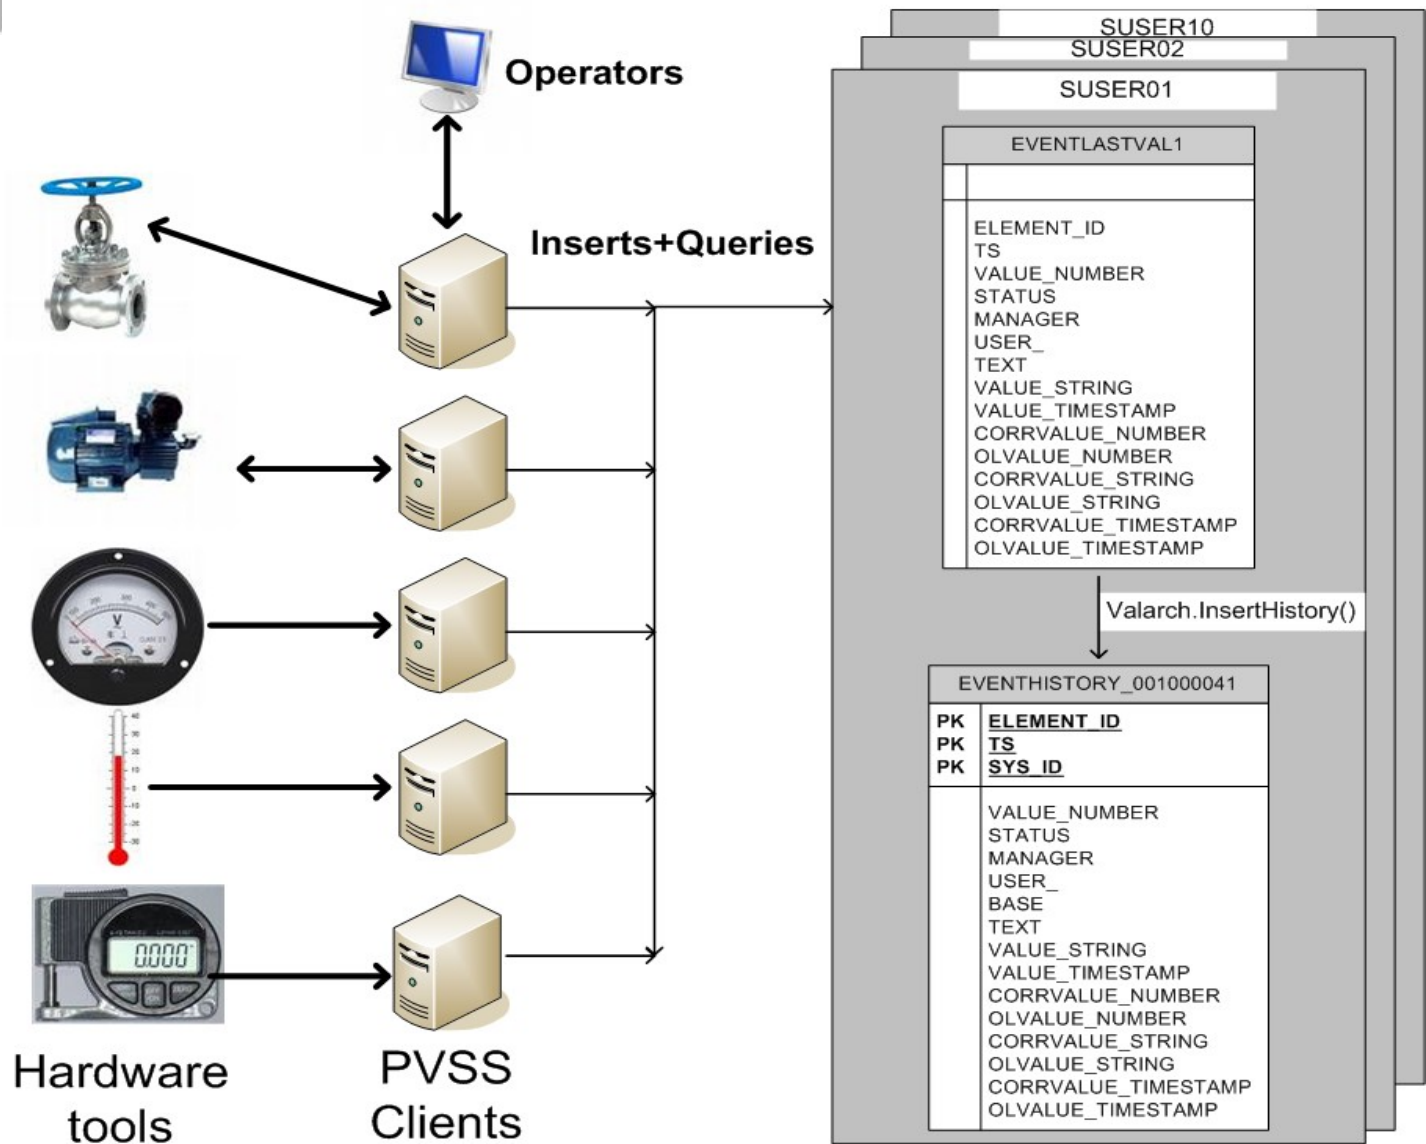
# Testing Exadata Storage ( 2 )

- 2<sup>nd</sup> phase of testing, 23-24 Oct
- Exadata test setup in Reading, UK
- PVSS swingbench benchmark
- Continuation of the tests from August



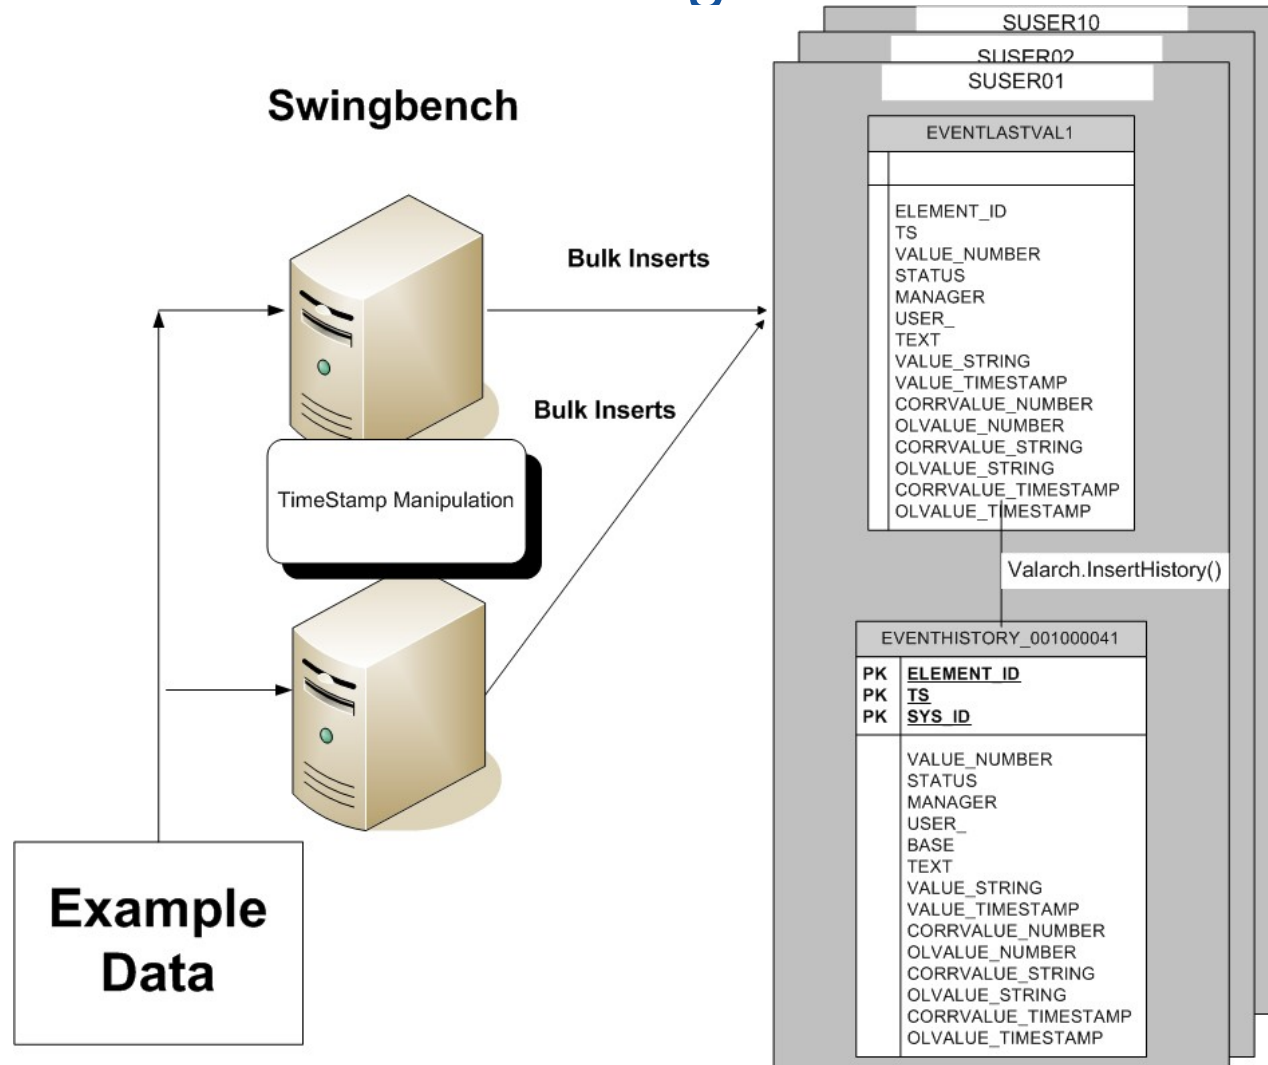
CERN  
openlab

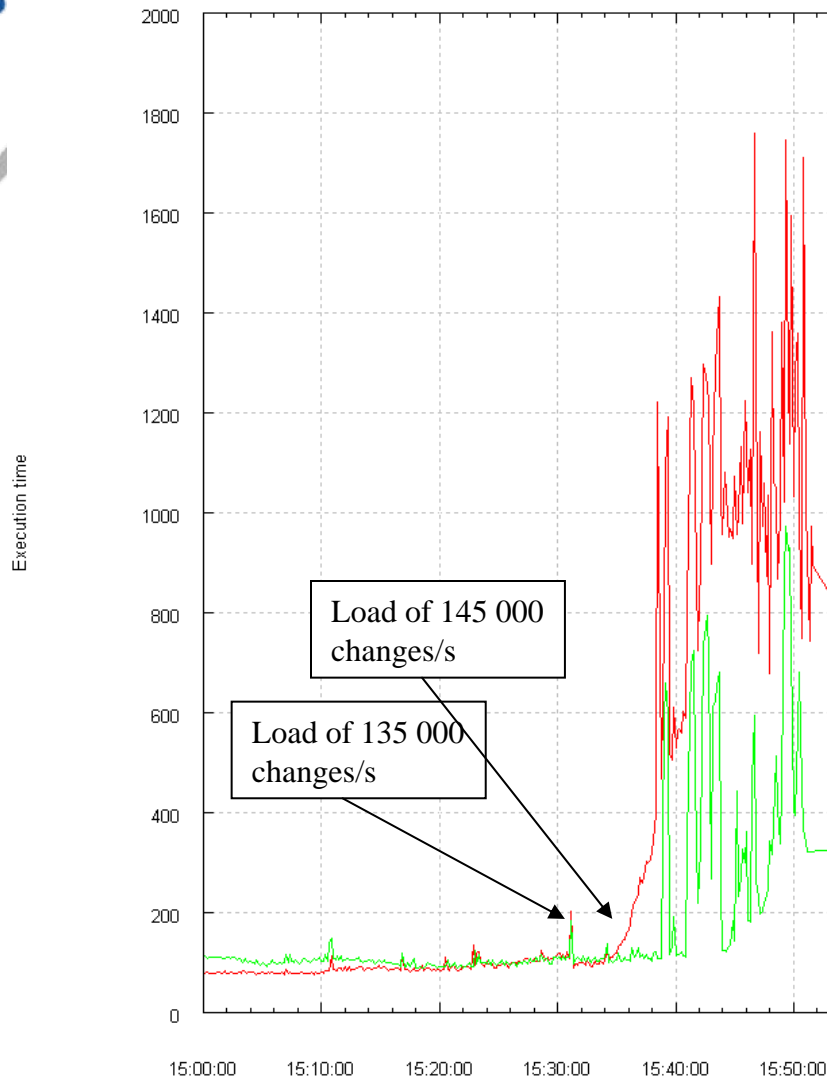
# PVSS Real Workload





- Created within Swingbench Framework

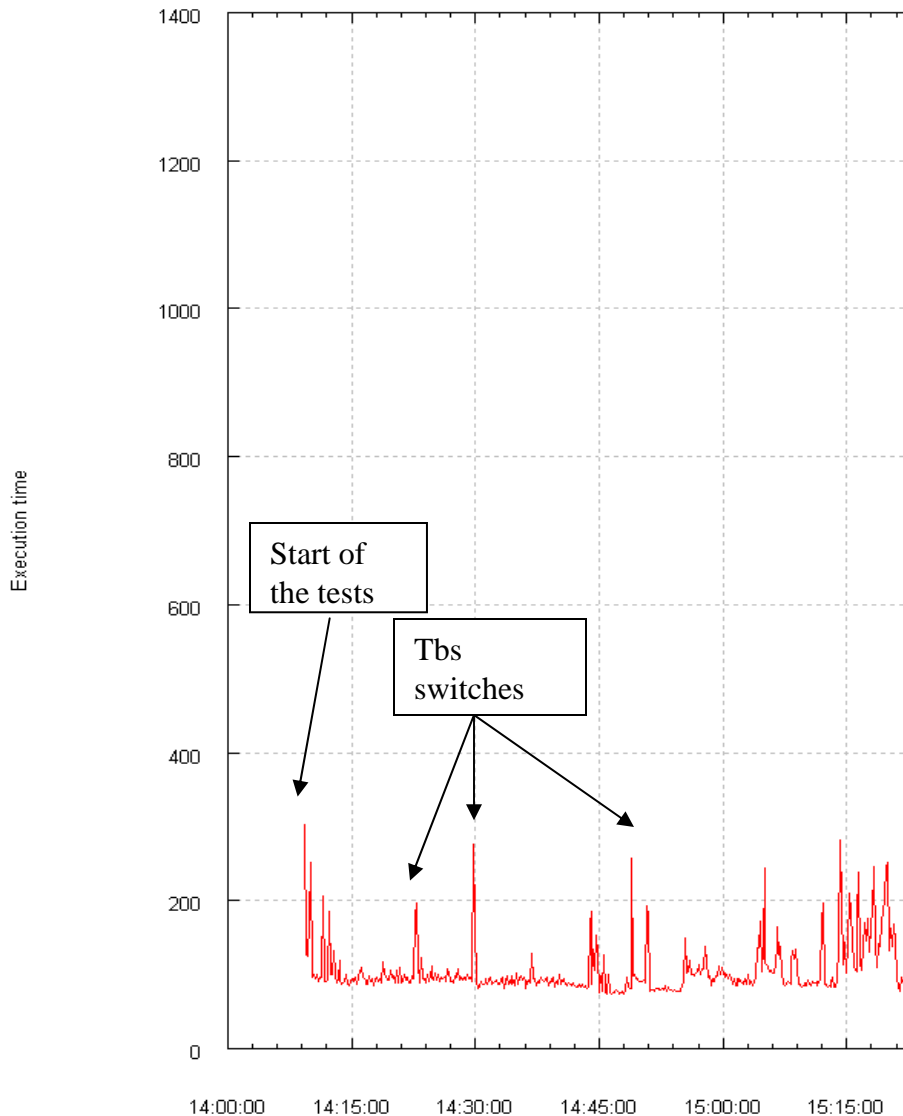




- Continued the tests started in August
- 4-Node RAC with 4 Cells storage
- 10 GB SGA
- 20 GB Tablespaces, 5 MB Uniform size
- Last stable point 145000 changes/s



# Fast file creation (1)




- Feature of Exadata
- `_cell_fcre = true`
- 17 seconds for 20 GB Tablespace
- Little spike in execution times
- Much below 1000 ms
- threshold



# Fast file creation (2)

Execution time



- `_cell_fcre = false`
- ~ 2 minutes to create 20GB tablespace
- Much bigger spike
- Higher than 1000ms threshold

- 4-Node RAC setup with Exadata storage:
  - could sustain up to 145 000 changes/s
  - bottleneck on concurrent change of control files
  - Much faster file creations lead to less spikes in execution times
- Overall better performance with Exadata storage features on.
- Next steps:
  - possibility get hardware onsite
  - or get Exadata software installed on our storage

# Q & A