Exadata, the new Oracle database machine

7 October 2008

Anton Topurov
Eric Grancher
Elzbieta Gajewska-Dendek

Outline

- Oracle Database Machine
- Exadata Storage Server Hardware
- Exadata Architecture and Features
- Beta Testing
- Test Results
- 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 Hardware

Exadata Storage Server

Racked Exadata Storage Servers

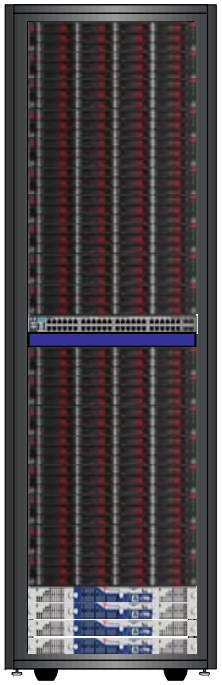

- **Building block of massively parallel Exadata Storage Grid**
  - Up to 1GB/sec data bandwidth per cell
- HP DL180 G5
  - 2 Intel quad-core processors
  - 8GB RAM
  - Dual-port 4X DDR InfiniBand card
  - 12 SAS or SATA disks
- Software pre-installed
  - Oracle Exadata Storage Server Software
  - Oracle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

Exadata Important Features

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

Exadata Architecture I

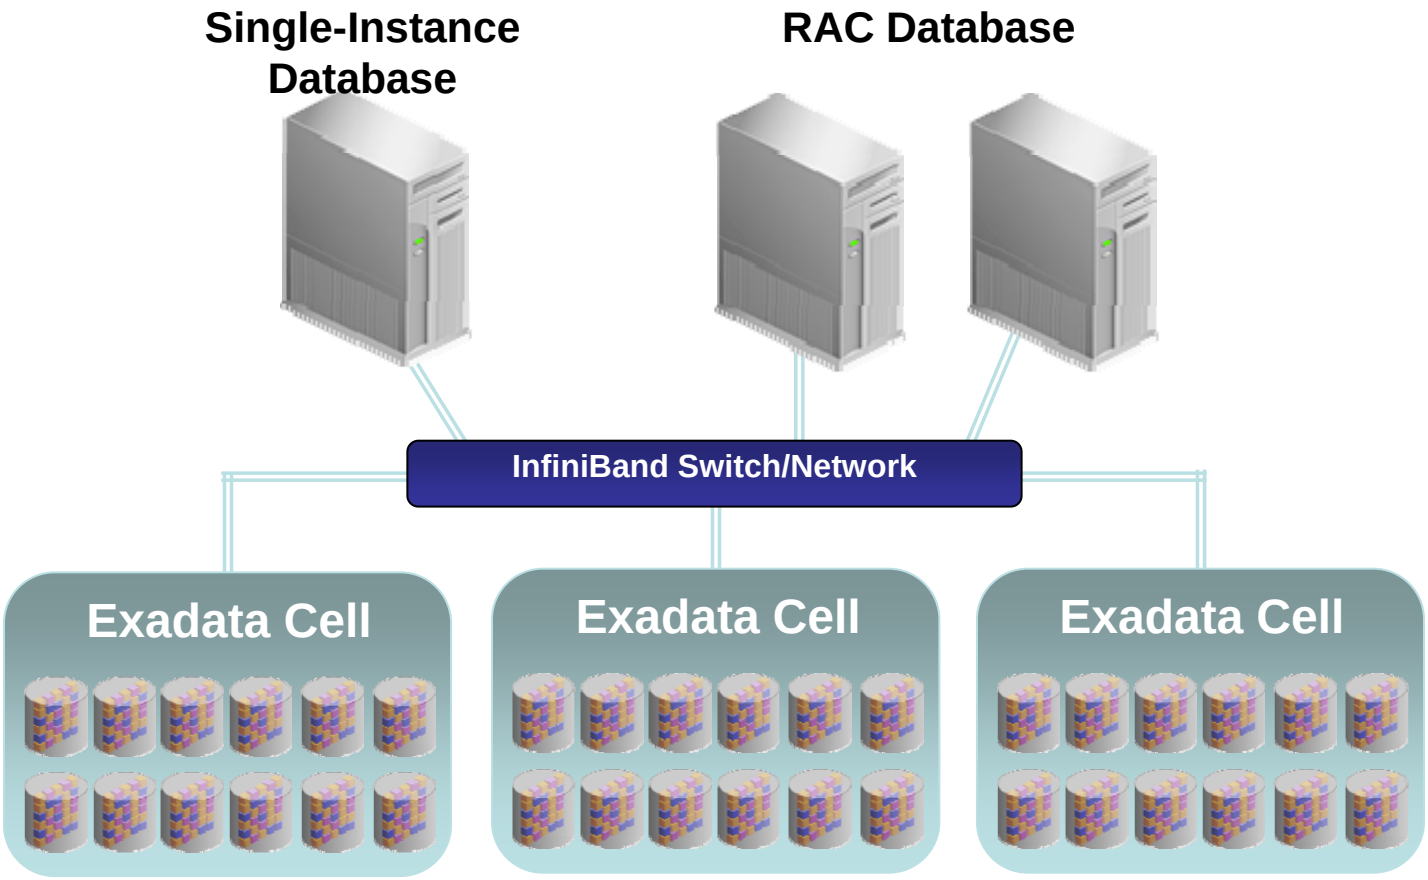

Exadata Architecture II

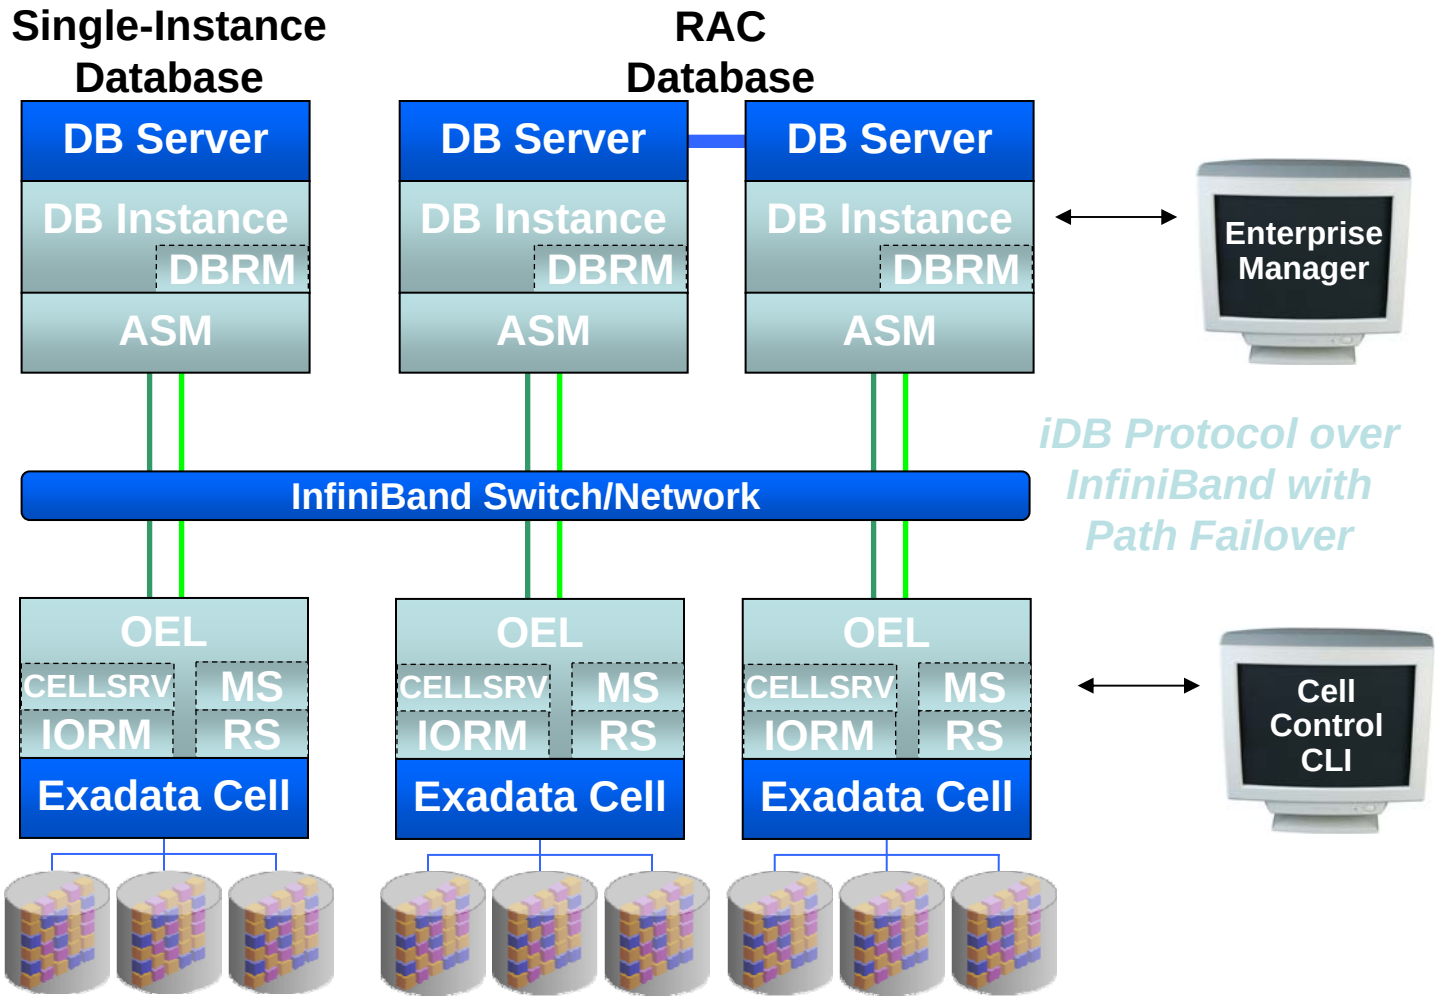

Testing Exadata Storage

[JST]: Exadata beta testing

- Beta testing since April
- 3 day learning event at Oracle UK site
- Developed PVSS swingbench benchmark
- Remote tests done in August
- More remote tests to be done this month
- Possibility to get the hardware onsite

PVSS Real Workload

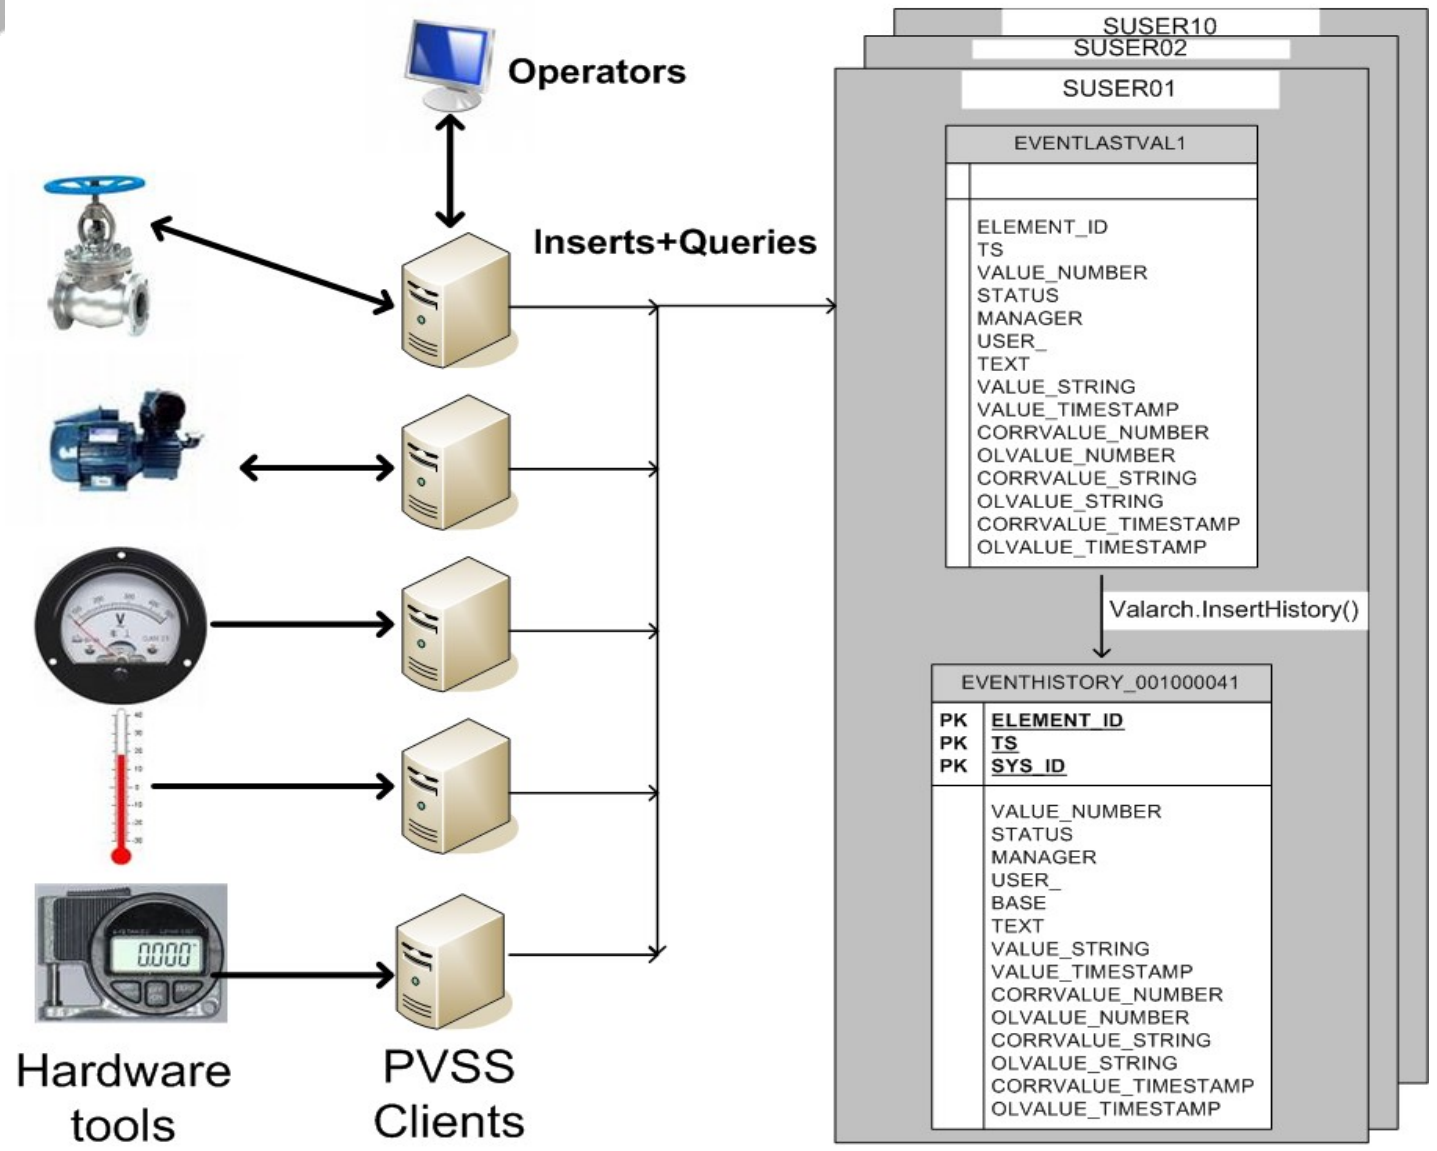

PVSS Workload Simulator

- Created within Swingbench Framework

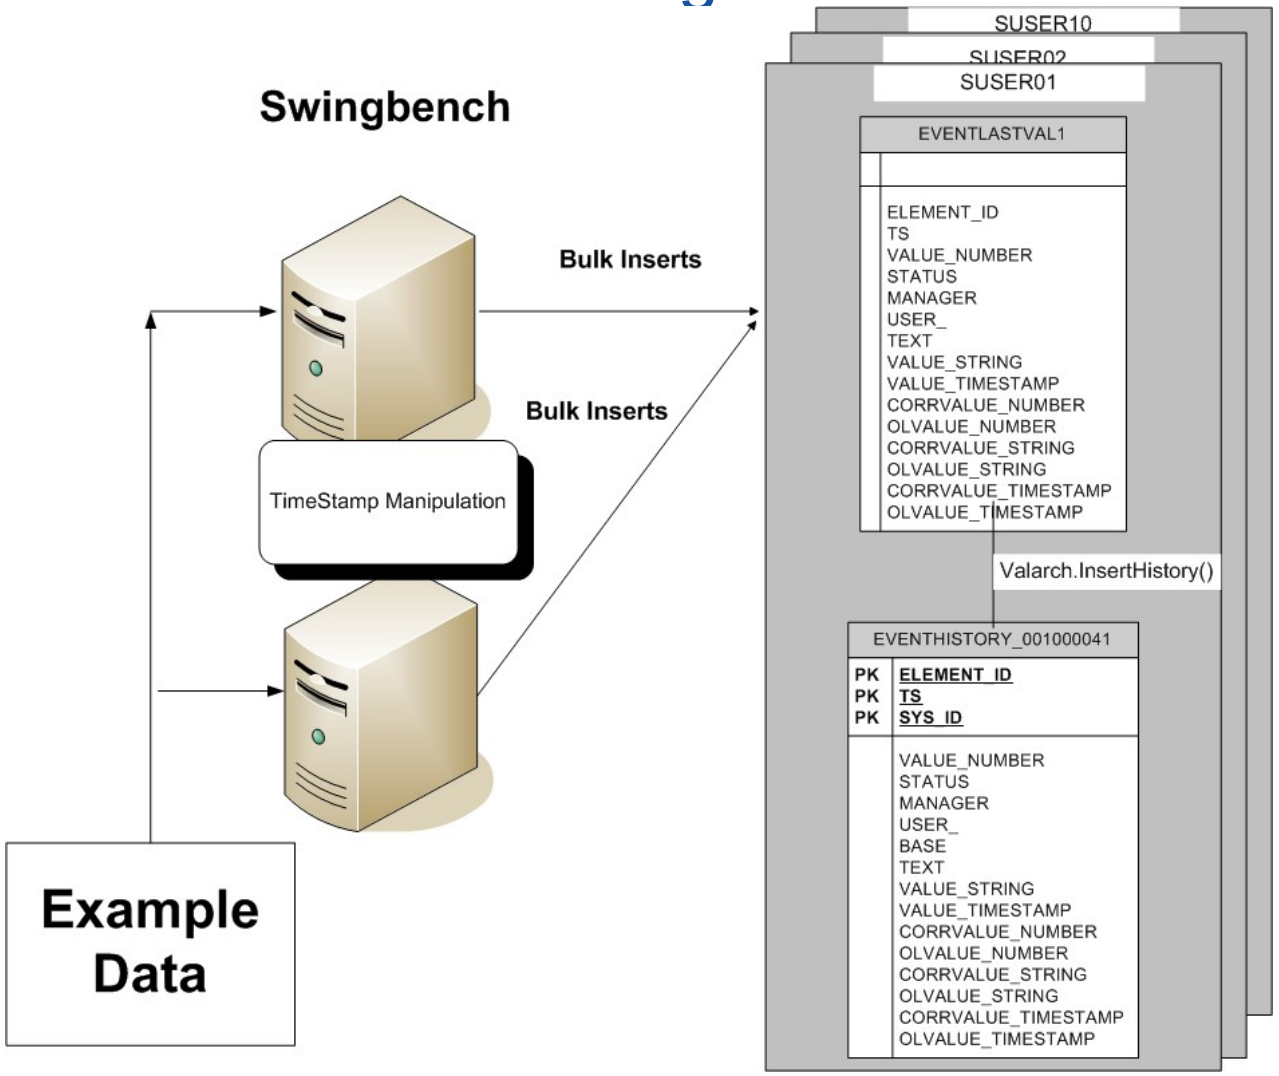

Swingbench in action

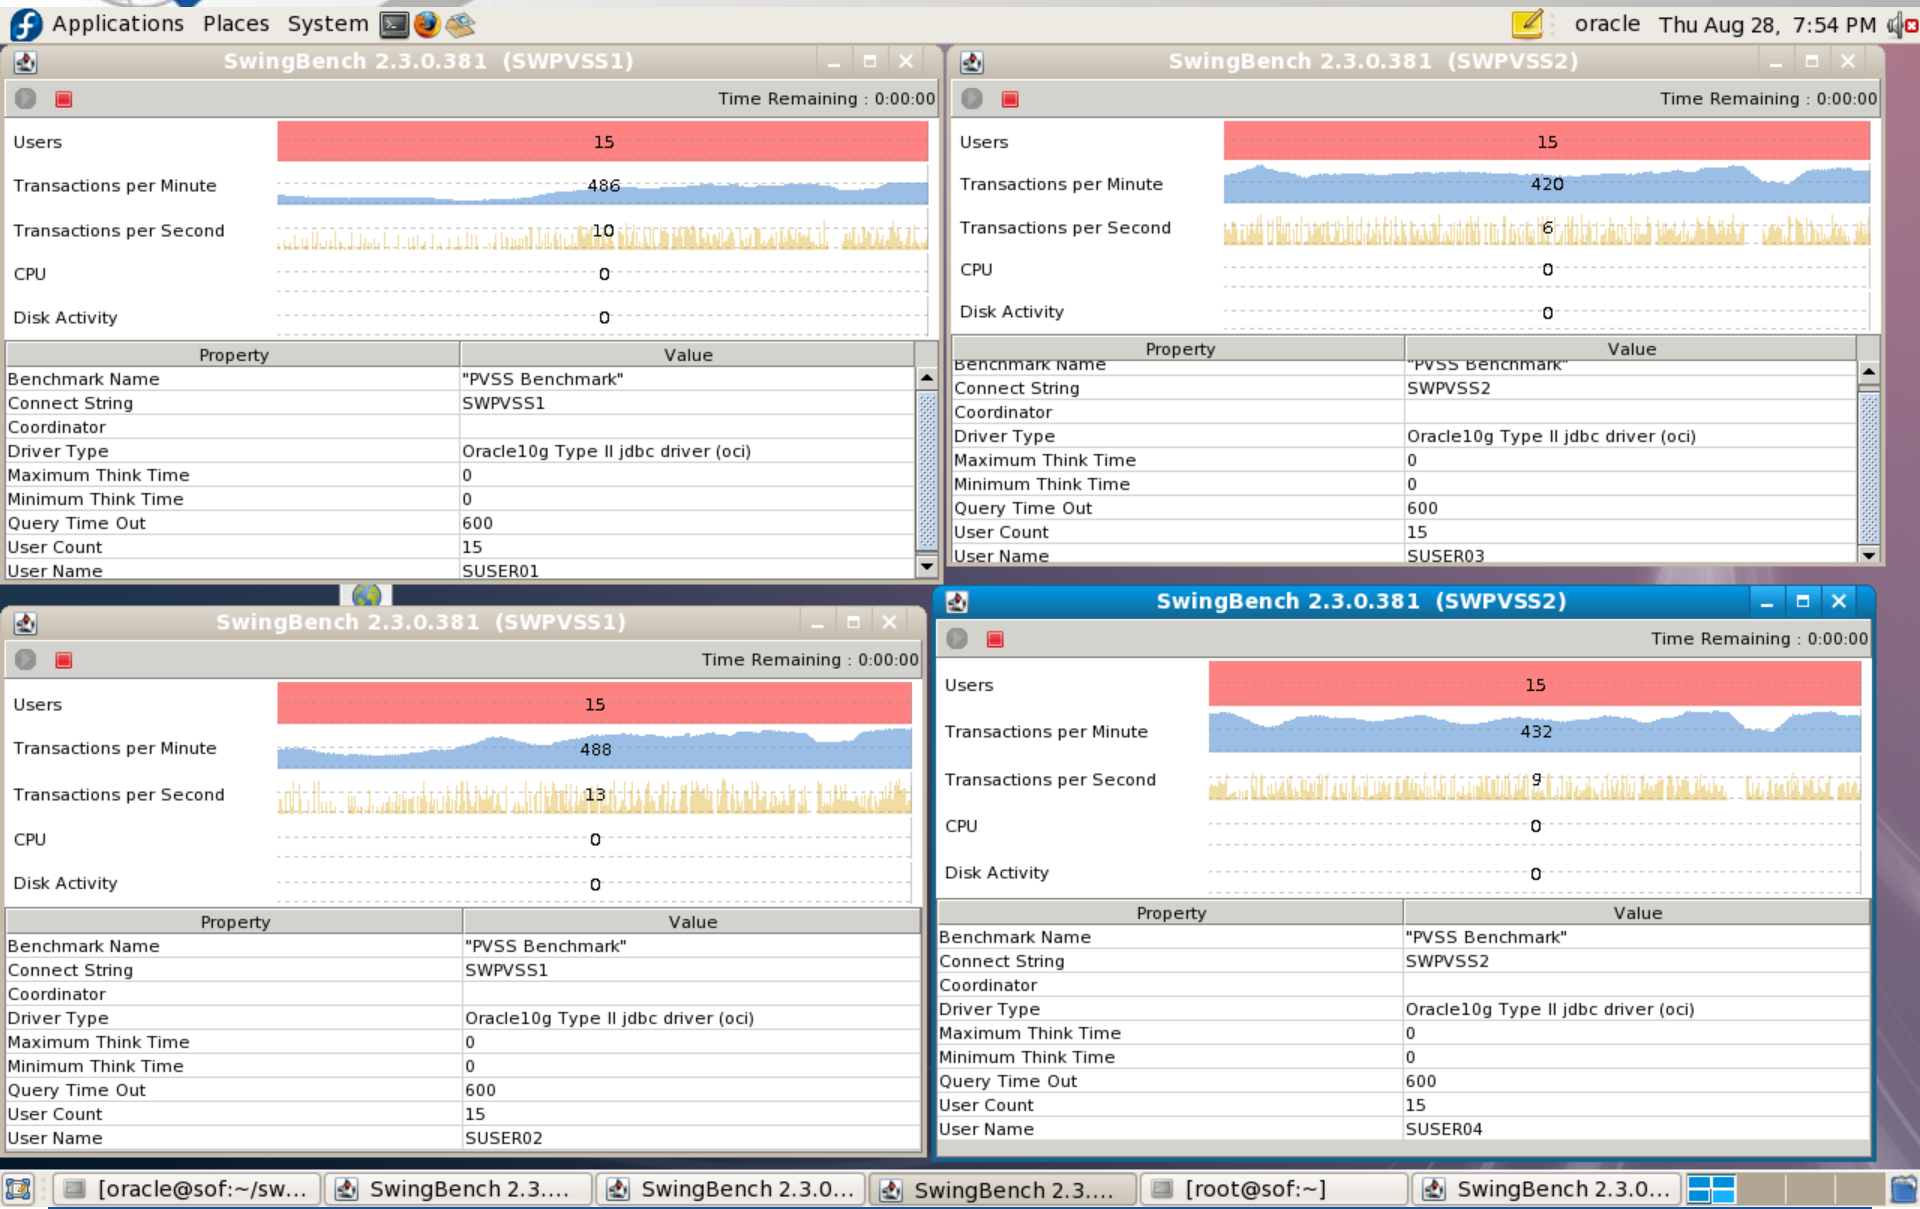

CERN setup test results

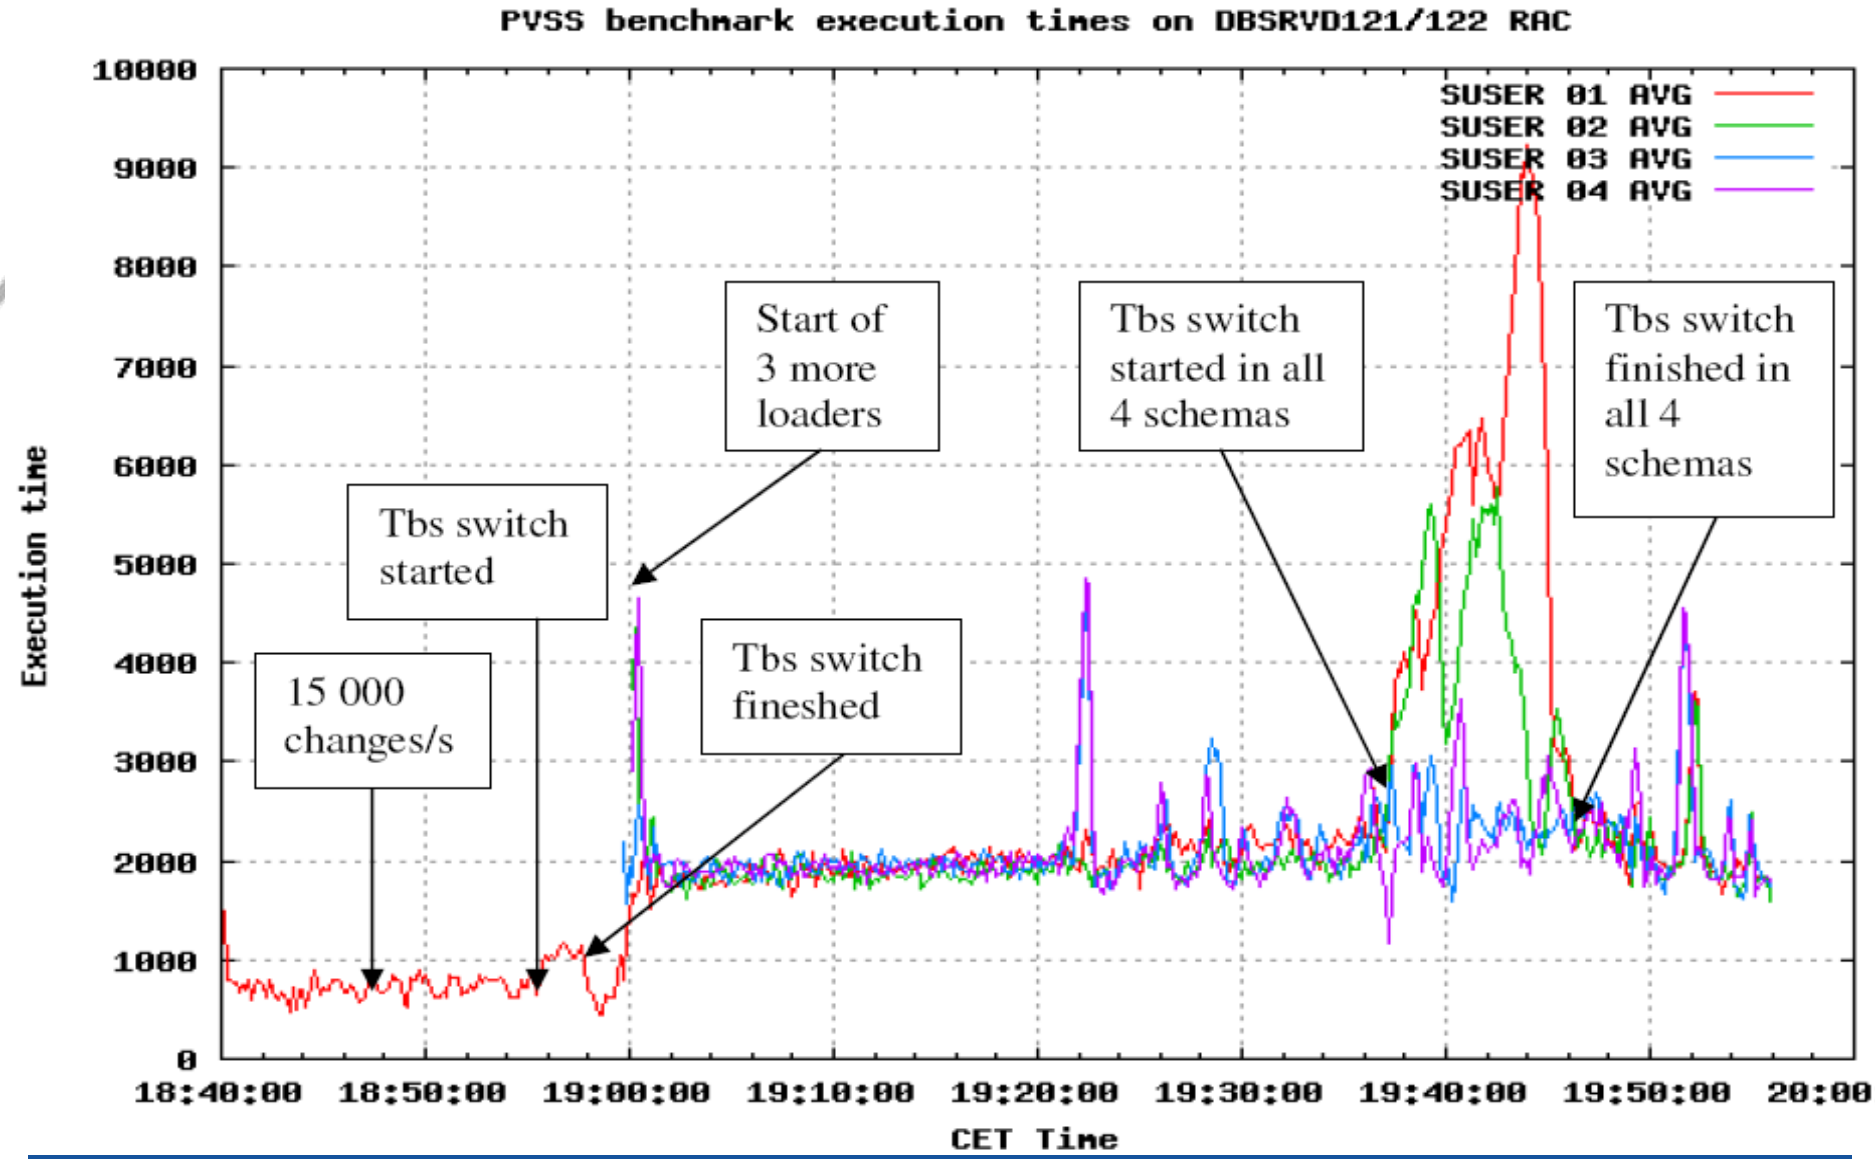

Exadata test results

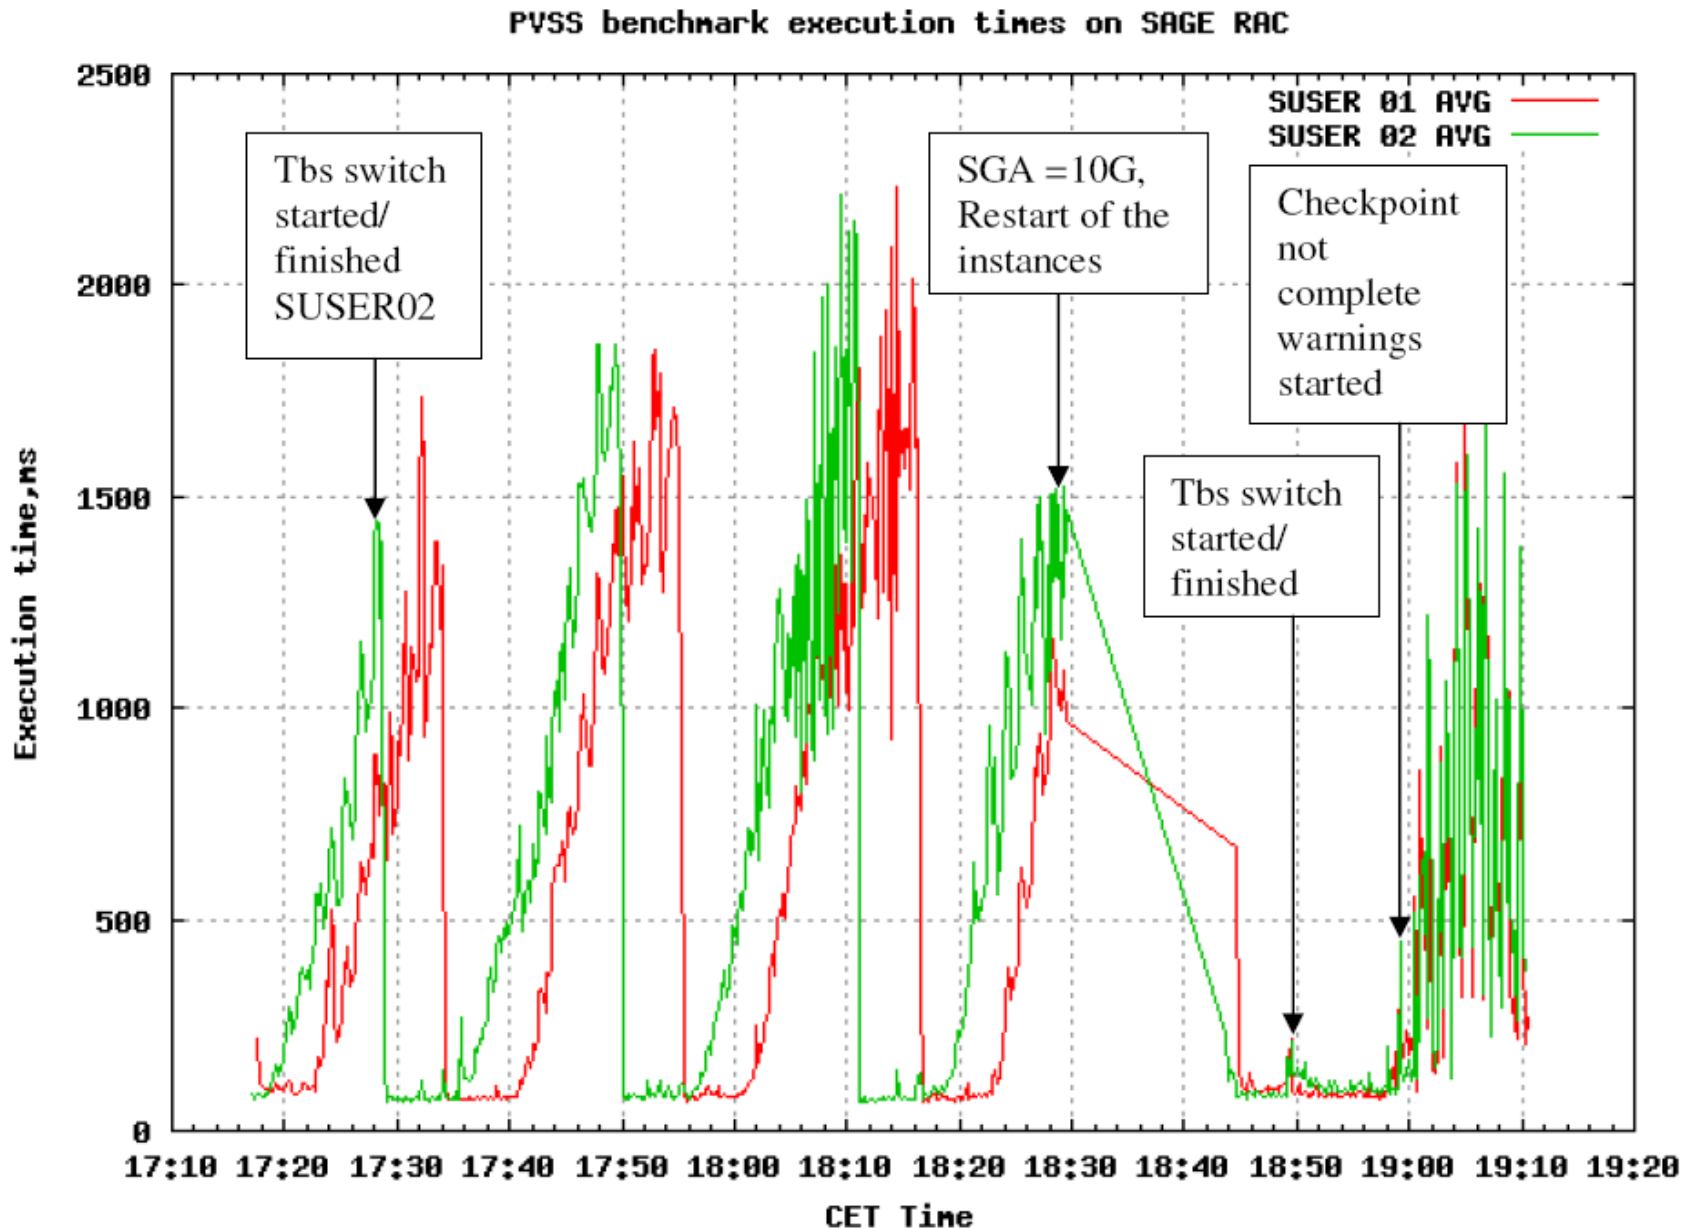

Conclusions

- **CERN setup**
  - 2 node RAC can sustain 30-40 000 changes/s
  - Degradation of performance due to tablespace creation
- **Exadata setup**
  - We have only intermediate results, more testing is needed
  - 4 node RAC, could cope with 150 000 changes/s
  - Performance degradation due to small SGA (1G)
  - With sga_target=10G, much better performance
  - Redo logs were the bottleneck
  - More tests scheduled to identify the full performance of Exadata

Q & A