

## CERN Database Services Review

Jacek Wojcieszuk, CERN/IT-DM

Distributed Database Operations Workshop


April 20<sup>th</sup>, 2009




- Current status of Physics Database Services
- Plans for 2009/2010 run

- Computer Center:
  - ~ 25 RAC databases (up to 4 nodes)
  - 114 servers, 164 disk arrays (almost 2000 disks)
  - ~550 CPU cores, ~ 900GB of RAM, ~750 TB of raw disk space, ~190 TB of effective disk space
- Took over the responsibility for supporting the online databases
  - **Additional funding** from the LHC experiments for the DBA team
  - Only ATLAS still missing
- Online databases (at Point 2, 5, 8)
  - 3 RAC databases (up to 6 nodes)
  - 16 servers, 18 disk arrays (~250 disks)
  - 52 CPU cores, 136 GB of RAM, ~100 TB of raw disk space
- Team of 6 DBAs + service coordinator and link to experiments
- Coverage **outside working hours** on “best effort” for production databases

- Computer Center:



- ALICE, CMS and LHCb on-line databases:

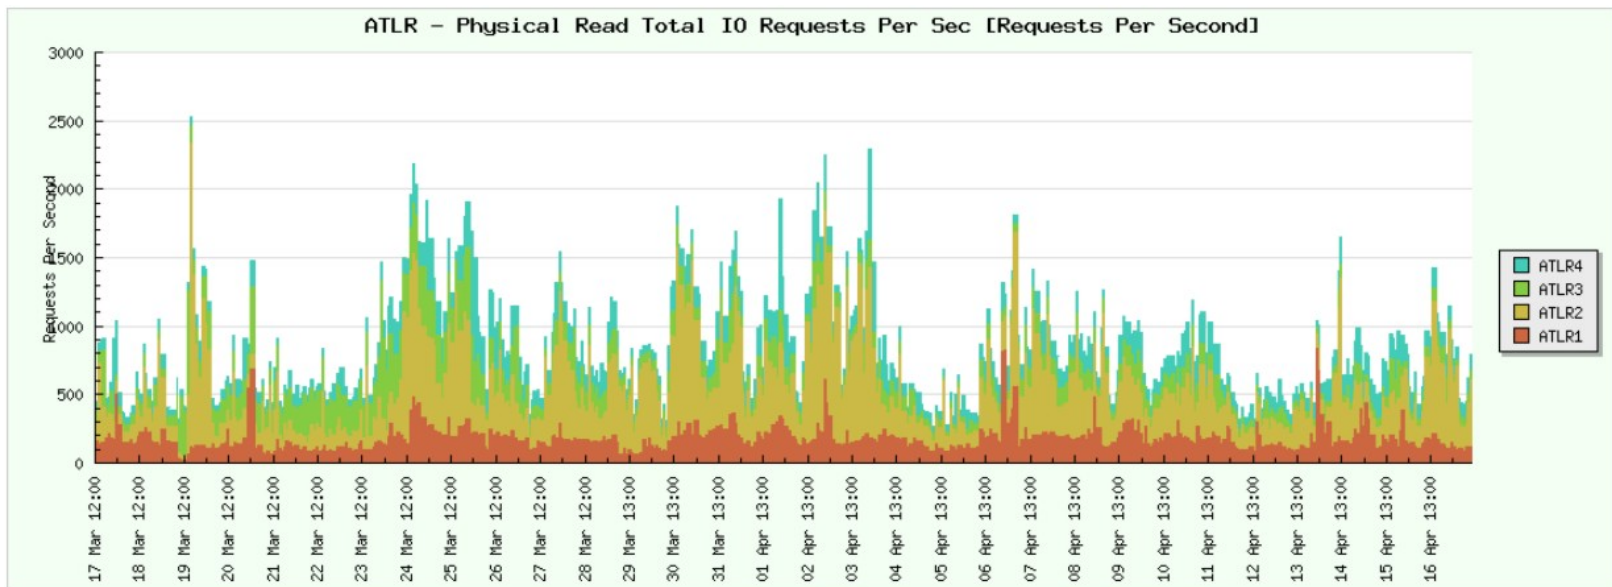



Database	Current Size	# servers	# arrays	Max Data Volume	Standby during 2008 run
ALIONR	0.8 TB	6 (SC)	3 + 1	11 TB	Y
ATLR	5.1 TB	4 (QC)	7 + 2	9.1 TB	Y
ATONR	2.4 TB	3 (QC)	4 + 1	5.2 TB	Y
CMSR	1.5 TB	4 (QC)	7 + 2	9.1 TB	Y
CMSONR	1.4 TB	6 (DC)	8 + 2	6 TB	Y
LHCBR	0.6 TB	3 (QC)	5 + 2	6.5 TB	Y
LHCBONR	0.5 TB	4 (QC)	2 + 2	3.5 TB	Y
LCGR	3.4 TB	4 (QC)	5 + 2	6.5 TB	Y
COMPR	7.2 TB	3 (QC)	8	10.4 TB	N


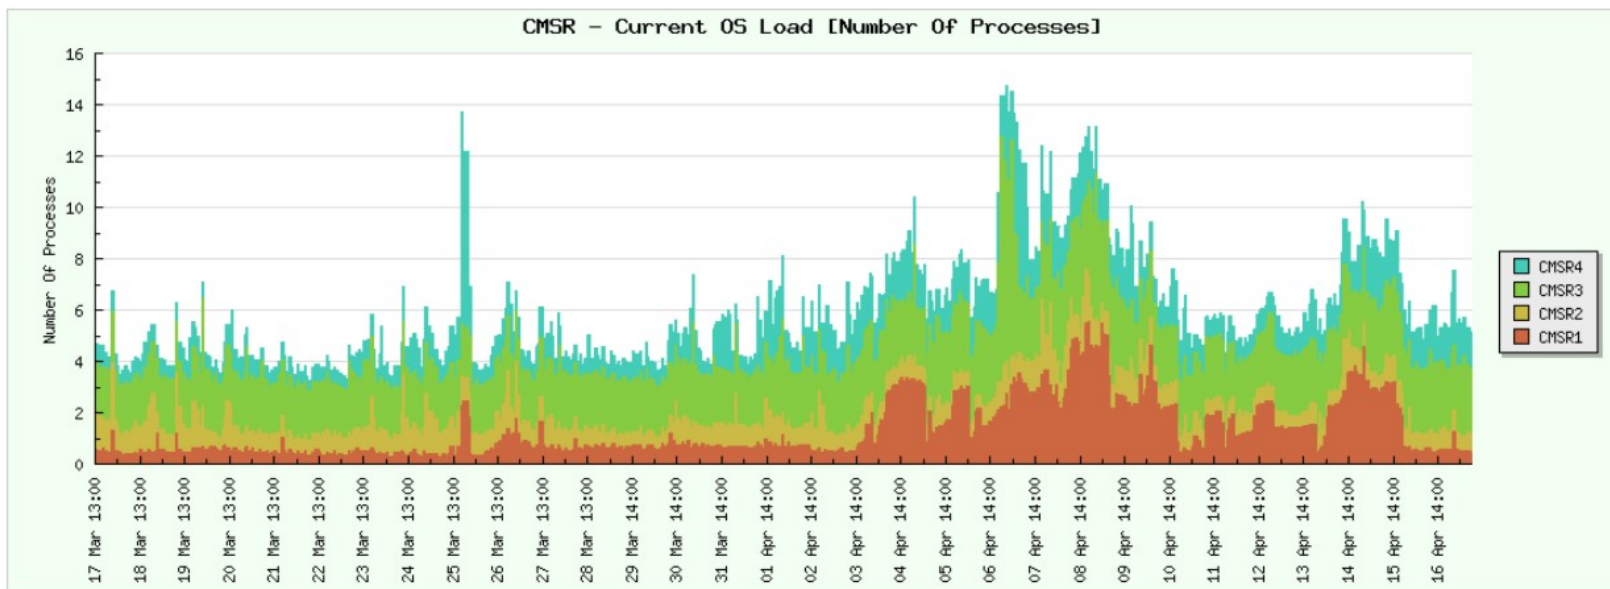
- All integration/pre-production run on the old hardware that will be replaced in the end of 2009
  - 2 x single core Xeon CPU servers
  - 8-bay disk arrays
  - will be replaced in the end of 2009



Database	Current Size	# servers	# arrays	Max Data Volume
ATLAS	1.5 TB	4	6	2.6 TB
CMS	1.1 TB	4	6	2.6 TB
LHCB	0.4 TB	2	3	1.3 TB
WLCG	0.5 TB	4	6	2.6 TB

- All the production databases seem to be sized correctly
  - The average load does not exceed 50% of hardware capacity
  - Saturation of certain resources seen occasionally only
- Data volume not an issue (for the time being)
  - Physical disk capacity still increases faster
  - ASM scalability (proven in tests) gives us confidence









- The service is ready for LHC start-up
- Required major changes implemented before **July 2008**
  - Production databases migrated to dual Quad Core servers
  - 64bit version of RHEL 4 and Oracle 10.2.0.4
  - Mixed power configuration: **half** of servers and **all** disk arrays connected to the critical power (power line protected with diesels)

## Plans for 2009/2010 run

- Replacement of hardware going out of warranty
  - Now: RAC7 (replacing RAC2)
  - In the end of 2009: replacement of RAC3 and RAC4
- Data Guard Physical Standby Databases
- Database 10.2.0.5 patchset?
- Upgrade of Grid Control to version 10.2.0.5
- Migration to RHEL 5.0
- Backup&recovery improvements
- See my next talk for details

Database	# servers	# arrays	Max Data Volume	Standby ?
ALIONR	6 (SC)	3 + 1	11 TB	Y
ATLR	5 (QC)	17 + 6	23 TB	N
ATONR	3 (QC)	5 + 2	5 TB	N
CMSR	4 (QC)	6 + 2	6 TB	Y
CMSONR	6 (DC)	8 + 2	6 TB	Y
LHCBR	3 (QC)	5 + 1	2 TB	Y
LHCBONR	4 (QC)	2 + 2	3.5 TB	Y
LCGR	4 (QC)	5 + 2	7 TB	Y
COMPR	3 (QC)	8 + 3	10 TB	N

- CMSR, LHCBR, LCGR will be moved to RAC7



# Thank you

