

# Munich Tier-2/3 @ RZG and LRZ

Stefan Kluth

Max-Planck-Institut für Physik

SC4/Tier-2 workshop LMU

19.09.2006

# Introduction



Ludwigs-Maximilian  
Universität München  
Prof. D. Schaile  
ATLAS Group



Max-Planck-Institut für Physik  
München  
Prof. S. Bethke  
ATLAS Group



Leibniz Rechenzentrum  
in Garching



Rechenzentrum der MPG  
Garching

# Introduction

- LMU/MPI/LRZ/RZG Collaboration
  - Build one nominal ATLAS Tier-2
  - additional 50% of resources as Tier-3
  - extra resources for MDT calibration
- Support
  - ATLAS Tier-2 and local Tier-3 computing
  - ATLAS MDT calibration centre
- Federation of two locations
  - Science campus Garching

# Hardware, 1/2 for ATLAS Tier-2/3

5 IBM BladeCenter chassis



8 Chenbro 3HE chassis, P4 server,  
Gb/s LAN, 12 SATA disks 300(250)  
GB on 3ware PCI SATA RAID



70 HS20 server, 2 x 3.06 Ghz  
Xeon, 2(4) GB RAM, Gb/s LAN



hp ProCurve 3400cl 44 port  
Gb/s switch



Total: ca. 180 kSi2k CPU, 28 TB disk  
incl. Eurostore external RAID arrays

# Hardware Experiences

- More than 25% of SATA Disks failed
  - Maxtor MaxLine II 250 GB
  - eventually replaced by MaxLine III 300 GB
  - no trouble since ...
  - except 1 mainboard
- BladeCenter HS20 servers stable
- one LAN switch died in thunderstorm

# Software

- Suse SLES9 OS
  - no problems
- Sun Grid Engine (SGE) batch system
  - supports cluster sharing, stability under load?
- MR-AFS global filesystem
  - semi-transparent migration to tape system
  - unstable under heavy load
- dCache
  - on two file servers
  - testing now

# Software

- glite/LCG 3
  - on SLC3 servers in DMZ
  - interface to SGE/MR-AFS accomplished
- Involves a lot of TLC
  - check SFTs
  - discuss with [dech-support.fzk.de](mailto:dech-support.fzk.de)
- Plans
  - full integration with LRZ part
  - fully support Tier2 activities

## Integration of an AFS-based Sun Grid Engine site in a LCG grid

### Objective:

Integration of existing Linux Cluster with

- SUSE Linux
- Sun Grid Engine batch system
- AFS File System
- operated by a Computing Center in the EGEE/LCG infrastructure


### Problems:

- LCG supports only Torque/PBS, LSF, Condor
- The computing center does not change its standard tools
- Integration of two security frameworks
  - > GSI
  - > Kerberos

### Approach:

- use Globus2 SGE-jobmanager & infoprovider developed by The London e-Science Centre\*
- adapt it for AFS and LCG-2.6

### Computing center Garching



- Existing cluster
- New users added (poolaccounts)
- CE authorized as SGE client
- No need to deploy software in the WNs:
  - ~ middleware in AFS, relocatable tarball installation of LCG 2.6.0

- Can be remote location
- Secure connection
- Acts as LCG gateway to cluster
- Poolaccounts get KRBS tokens via keyfiles

### Extra fixes:

- JobManager needs "use filetest 'access';" (Globus does not know about ACLs)
- Fix for job\_script\_name in SGE jobmanager
- Force bash usage in SGE wrapper scripts, source grid\_env.sh
- Fix quoting in globus-job-manager.conf

### Conclusions:

- restriction to Scilinux prevents LCG dissemination
- presented approach circumvents this obstacle
- SGE-AFS clusters can be integrated within the LCG Grid

\* <http://www.lesc.ic.ac.uk/projects/SGE-LCG.html>

# LCG/SGE integration

Ariel Garcia and Alberto Baragatti did all the work!



# Hardware Plans

- Hardware on tender
  - 2 BladeCenter à 14 HS20 2 x 3.2 GHz Xeon
  - corresponds to ca. 80 kSi2k
  - 25 TB file servers
- Near Future: 2007 and 2008
  - depending on final LHC schedule
  - ca. 230 kSi2k CPU (4-5 BladeCenter)
  - ca. 140 TB file server (25 file server à 5 TB)
  - larger/more LAN switch(es)

# Hardware Developments

- Multi-core CPUs
  - need 2 GB RAM/core or multithreading athena
  - increase LAN bandwidth of servers
  - Opteron + Cell solutions (QS20 blades)?
- SAS/SATA Fileserver
  - connect many SATA-II disks via SAS expander to 3 Gb/s SAS channel on server
  - connect ca. 7 TB @ 50 MB/s/TB per channel
  - RAID in hardware or software?
  - cost around 1 k€ / TB possible

# Status LRZ/LMU part

- Integration into existing Computing centre setup & infrastructure:
  - Suse, SGE batch
  - started with parasitic use of existing resources
  - NorduGrid/ARC in 2004, move to LCG 2005/06
- Substantial funding received for this year still
  - 150 kSI2k, 40 TB
  - full integration into WLCG/EGEE in next week