

WLCG GridKa+T2s Workshop

Site Report

DESY-ZN

Andreas Haupt, Peter Wegner

DESY, Germany



Worldwide LHC Computing Grid
Distributed Production Environment for Physics data Processing



overview (1)

- DESY Tier2 centre for ATLAS and CMS
 - DESY-HH: offers $\frac{3}{4}$ of the resources
 - DESY-ZN: offers $\frac{1}{4}$ of the resources
(according MoU for Tier2s)
- Grid services offered by the DESY-ZN site
 - 2 CEs (Torque, SGE)
 - SRM-enabled SE (dCache)
 - GIIS, BDII, MON box
 - supported VOs: atlas, cms, dech + DESY VOs (ildg, zeus, ilc, hone, ...)



overview (2)

- OS / middleware
 - SL3 / 32bit (from Fermi)
 - gLite 3.0.2 (RPMs, tarball)
- Hardware status, Grid farm
 - 50 dual Xeon farm nodes (3.2 GHz, 2 GB memory, 50 GB scratch space)
 - 3 TB disk storage on SE
- Network connection
 - 155 Mbps, XWIN; 10 Gbps VPN planned



system installation / maintenance

- no real framework is used
 - installation and configuration just done by kickstart, rpms, yaim
 - kickstart post install
 - rpm postinstall scripts, triggers (also run yaim)
 - rpm updates via yum cronjob
 - rpm selection provided by self-written yumsel script
 - configuration files installed per rpm



storage element

- already existing dCache instance has been extended
 - dCache 1.6.6
 - SRM door (including gsidcap door)
 - disk pools servers (hosting also gridftp doors)
- problems:
 - interrupted gridftp transfers lead to files filling up the pool nodes but do not show up in pnfs → pools get cleaned regularly „by hand“ , will be fixed in dCache 1.7.0
 - dcap, gridftp do not honour pnfs directory permissions correctly
 - integrated information provider still not fully usable



computing elements (1)

- Torque CE
 - used torque, maui provided in the gLite repository
 - batch master, maui scheduler run on a separate node
 - avoids batch system reconfiguration on CE reinstallation
 - but problems with apel accounting
 - fair share policy accomplished with maui
 - vo software manager accounts get higher priority
 - vo production manager get higher fair share than „normal“ users



computing elements (2)

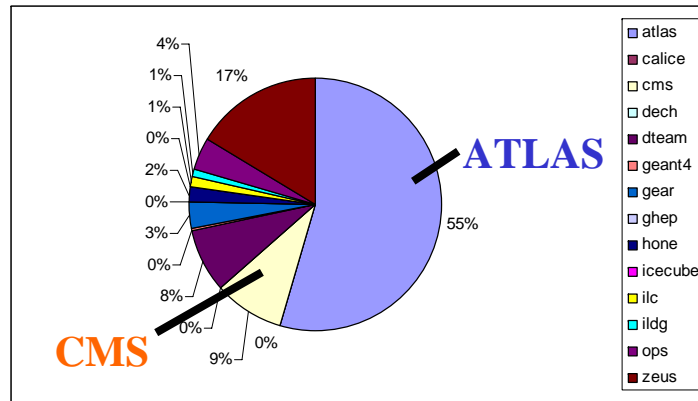
- SGE CE

- opens our internal farm to the LHC grid (internal farm: ~300 cores, 64-Bit OS, access to HSM, ...)
- only for testing purposes by now (only ops, dteam supported)
- gLite 3.0.2 lcg-CE rpms + adopted lcgsgc.pm job manager (original lcgpbs.pm)
- information provider is an adopted version of the one provided by Imperial College (London)
- farm nodes (64 bit SL3) use lcg 2.7.0 tarball installation (gLite 3 tarball was not functional - seems to be solved now)

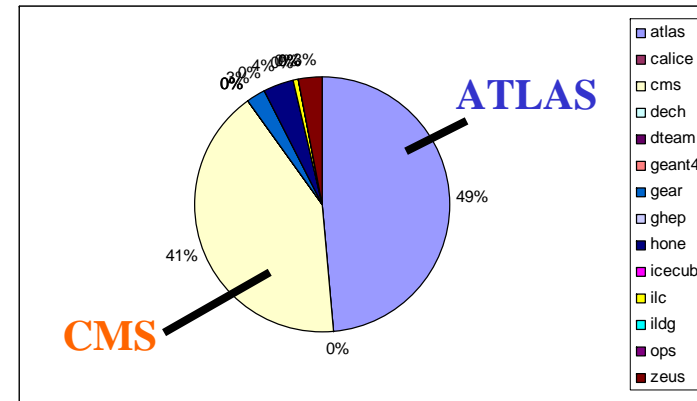


grid farm statistics

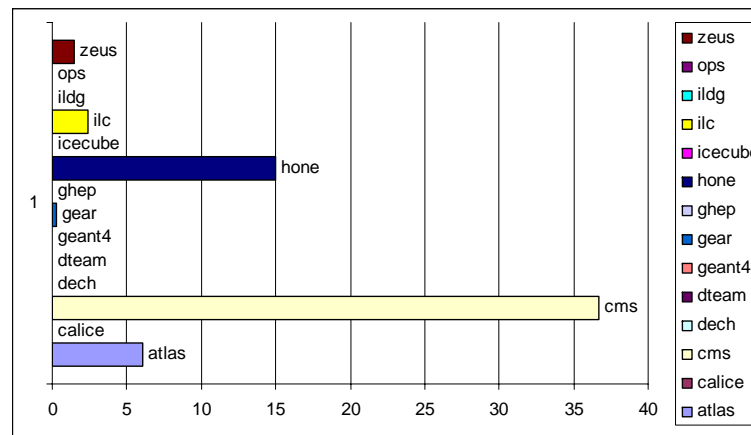
Number of Jobs



Wall clock time



CPU time per Job





SC4 participation

- took part in the atlas SC4 from June 2006
 - dq2 dataset subscriptions
 - results are rather disappointing (shocking)
 - very slow transfers in comparison to „normal“ srm / fts transfers
 - lots of errors (timeouts, lfc problems)
 - actually no complete dataset arrived us



Conclusions

- Grid in production since April 2006 for ATLAS, CMS and all DESY supported VOs, meanwhile stable
- Grid management (OS/MW) successful with simple tools
- Usage of both Torque and SGE (test) CEs
- More VOMS based security necessary for SEs and other services
- SC4-, FTS channel tests only partly successful