

# NAF status

... and some details on Lustre (courtesy Götz Waschk)

Andreas Haup, [Yves Kemp](#)

NUC August 2011 meeting

Place, Date

# Operations

- > Smooth running, many jobs
- > Upgrade to SL 5.6
- > Login troubles for ATLAS on 1.8. -> solved
- > Some users experienced troubles with tcx200 (48-core AMD test system), see later



# Downtime(s) since last NUC

## > 29.7. – 1.8.: Power shortage on campus and in CC

- Reduced NAF and Grid power in HH – all went well

## > Network interventions due to LHCone

- Did not affect users directly – most of them transparent
- Problems on KIT side resulted in GGUS against DESY-HH and DESY-ZN
- After some fiddling around, all OK now

## > À propos LHCone:

- DESY (HH+ZN) now part of HEPPI – dCache traffic to/from KIT goes via dedicated 10 GE link
- Link has been tested: Single stream iperf at around 1.8 Gbit – **OK**
- Aggregated BW peaks seen over 9 Gbit with iperf tests - **OK**



## Reminder: Downtime 25.8

- > As announced in the last NUC meeting, we will have a larger downtime on 25.8.
- > Preparations running – more information in due time
- > ... further downtime at Zeuthen site on 1.9. possible – not confirmed tough.
  - Would not be possible to combine with 25.8. downtime
  - Will try to keep overall impact smallest possible
  - Planning once confirmation is there



# Batch system usage

## > MultiCore and Single-Slot usage since 1.7.2011

VO	MultiCore (wall-day)	Single-Slot (wall-day)
atlas	454	10475
cms	12	39777
lhcb	0	4149
ilc	0	2625

- > Sum: 57026 Wall-Days since 1.7.2011 (all jobs)
- > 1564 slots available now – 40 days since 1.7.2011
- > 62560 max Wall-Days since 1.7.2011 (in reality somewhat less ...)
- > **91% occupancy**



# NAF additions

- > AMD and Intel will bring out new systems in next months
  - E.g. WLCG contribution from DELL  
<https://indico.desy.de/getFile.py/access?contribId=50&resId=0&materialId=slides&confId=4019>
  - Especially AMD Interlagos very interesting platform
- > No benchmarks (synthetic, HEPSPEC06 and real applications) available yet – but we are in contact with KIT who will have test systems
- > AMD Interlagos systems will most likely be systems with 4 cpu a 12 cores. How do such systems behave in Grid and NAF?
  - **Grid-HH:** Have such a system (current CPU generation, 48 cores, 196 GB RAM, 6 Disks). Since half a year: Good performance, no operational issues, no bottlenecks, no noticeable performance degradation (according to CPU/WALL metric)
  - **NAF:** Tested such a box: tcx200: CPU/WALL efficiency constant for some users, dropping for some users, and being catastrophic for other users during the two days of the test. Also seeing one SGE problem. **Need to understand this!**
- > Plan: Know what to buy by beginning October ... buy October/November



- > 7 Lustre file systems, 575 TB total
- > 21 server machines
- > Running Lustre versions 1.8.4 or 1.8.5 on servers, 1.8.5 on clients
- > Quota active everywhere, used only for daily status reports on usage
- > Networking
  - Infiniband and bonded 1G Ethernet used in Zeuthen
  - HH: After 25.8.: Server 10 GE, Clients (mostly) 1 Gbit and (some) 10 GE, Infiniband to be removed

- > Directly attached storage only, on Dell und HP JBODs
- > File system /scratch/zn/lustre/lhcb
  - Servers: tcmds1, tcl01, tcl02, tcl03, tcl04
  - Used: 13/47 TB
  - End of warranty: 2011-08-21 (tcmds1)
- > File system /scratch/zn/lustre/atlas
  - Servers: tcmds2, tcl05, tcl06, tcl07
  - Used: 18/20 TB
  - Warranty: ended
- > File systems /scratch/zn/lustre/atlas2/{monthly,users}
  - Servers: tcmds2, tcl08, tcl09
  - Used: 29/30 TB, 0.7/6TB
  - Warranty: ended (tcmds2)



- > Mix of directly attached storage and SAN on DDN
  - DAS servers: tcoss107, tcoss108, tcoss109
- > File system /scratch/hh/lustre/cms
  - Servers: tcmgs12, tcoss105, tcoss106, tcoss110, tcoss111, tcoss112
  - Usage: 53/58 TB
  - End of warranty: 2012-01-20 (tcmgs12)
- > File System /scratch/hh/lustre/atlas
  - Servers: tcmgs12, tcoss105, tcoss106, tcoss107, tcoss108, tcoss109, tcoss110, tcoss111, tcoss112, tcoss116
  - Usage: 94/105 TB
  - Warranty: Ended (tcoss107, tcoss108, tcoss109)
- > File System /scratch/hh/lustre/ilc
  - Servers: tcmgs12, tcoss105, tcoss106, tcoss107, tcoss110, tcoss111
  - Usage: 29/33 TB
  - Warranty: Ended (tcoss107)

