

NAF status report

Yves Kemp

NUC meeting, 9.6.2010

General running:

- > Overall stable running
- > Occasional problems with (mostly) ATLAS IN.
 - High load
 - Kswapd's taking ~100% CPU
 - Presenting many Lustre errors
 - Ssh to not work
- > One Batch-WN recovered when left alone (and running out of jobs), IN WGS have to be rebooted for fast recovery. Interesting behavior...
- > Unfortunately no real clue what is the problem behind, and how to solve it
 - (Probably) Kernel bug, (eventually) triggered by Lustre bug (?).
 - Lustre client 1.8.3 did not help, min_size memory tuning did not help
 - Managed to get more forensic out of the systems ... analyzing



General running 2:


- > CMS observed Lustre problems also:
 - Probably due to an overloaded OSS. Not clear whether CMS or another experiment on the shared OSS are cause of the problem
- > CMS users have complained about bad performance and errors reading from DESY-HH dCache
 - One problem: Pools still on 1-Gbit, and reaching this limit. 10-Gbit on the agenda in the next weeks
 - Other problem (somewhat connected): Some users overload system with too many requests (4000 jobs on only a couple of files). Unfortunately, batch scheduler and storage system ignore each other, so this leads to overloading.
- > ATLAS and DESY-HH dCache
 - Generally fewer reading than CMS, so less problems
 - Pools also on 1-Gbit, will be upgraded, limit less painful as in the CMS case
 - Some user also reported problems: Writing to dCache w/o DQ2, one job opening too many files and blocking dCache
 - Little reading compared to writing (1W/2R Zeuthen, 1W/4R Hamburg)
- > dCache: We start to see **real analysis work!**




SL4 WGS

- > Very small load
- > List of users given to experiments
- > Still needed?


tcx030 CPU Load Last 48 Days



tcx031 CPU Load Last 48 Days



tcx040 CPU Load Last 48 Days



Lustre space & extension

> Extension

- HH: ~4 TB already assigned to ATLAS two weeks ago
- HH: Three new server with each 4x8 TB RAW (net ~26-29 TB) in preparation phase
- Will be used for the 20 TB request by CMS and 20 TB by ILC
- ZN: 16 TB added to LHCb recently

> General extension constraints:

- Known instabilities...
- For larger additions: Some boundary conditions need to be clarified.

> Scratch deletion

- No real news...
- ... except that one good idea proved to not work....

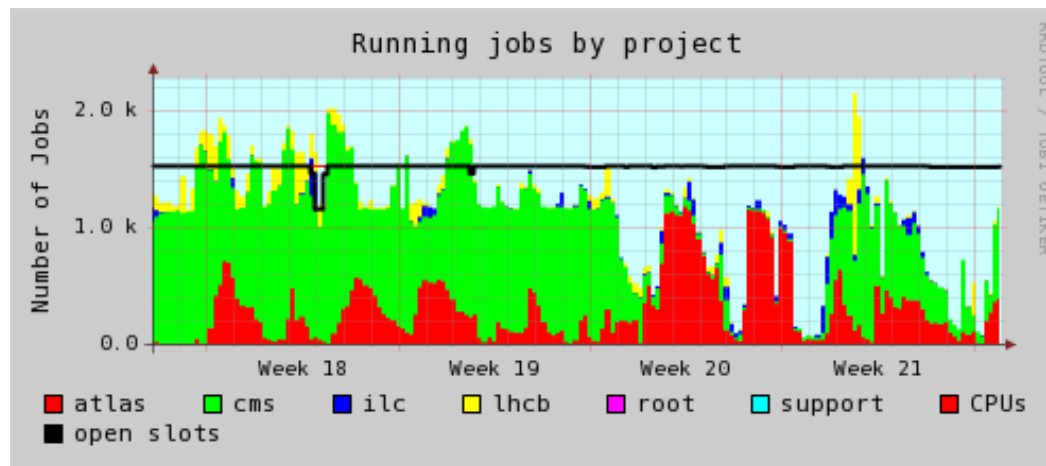


Monitoring & Accounting

- Multicore monitoring: below the 1%:

	1-Slot CPU-Days	N-Slot CPU-Days
ATLAS	13336	173
CMS	24970	366
ILC	1198	0
LHCB	1953	0

- Jobs last month:



- Generally: As a more long-term project, we will rework the accounting, also to join Grid (both DESY-HH and DESY-ZN) and NAF accounting.

- > PRC recommendation concerning NAF:
 - “The NAF is strongly used and very well accepted by the German particle physics community. However, the PRC is worried that the NAF might saturate in the future with increasing computing requirements from the LHC. For the next meeting, the PRC would like to see estimates for the future resource-requirements and an evaluation of the current load from different use-cases, like Monte Carlo production and data analysis.”
- > Next PRC is mid October 2010. We will need your help!
- > Unified accounting will already help (will be in place before this date)

NAF Helpdesk Tickets

Software Request `tkdiff` instead of `kompare`, JEM for Ganga,
`ctags`

Login Problem

- > No WGS available 5
- > Unspecified problem
- > MyProxy error message
- > Atlas User login
- > Strange warnings
- > WGS alias vanished

storage

- > corrupted files on Lustre
- > dCache hanging (DoS)
- > slow access to `afs/lustre 2`
- > no rights in home dir (temp)
- > file not available in dCache (temp)



NAF Helpdesk Tickets

Registration Not registered before first login

Batch system

- > Slow starting jobs
- > Accounting for second account
- > jobs dying - not enough cpu time requested

X11 forwarding

- > SL4 X11 forwarding did not work
- > Eqw state

Other .Oldfiles vanished

Summary Helpdesk Tickets

- > Helpdesk received 25
- > Forwarded to experts 15