

NAF Status Report.

Andreas Haupt

[Yves Kemp](#)

NUC phone conference

12.08.2009

Problems last month

- > No further crashes of the login VM hosts
 - Could not be provoked either
 - Machines running a newer kernel by now
 - Problem solved?
- > Some memory bars exchanged (as usual)
 - Machines with **unusual** RAM troubles should be migrated to SL5 to ease debugging
- > Linecard of central switch disturbed
 - see later for more details
- > DESY-HH: Half of CPU resources migrated to SL5
 - Grid-ce5.desy.de
- > Switch partial downtime 6.8. successful
 - Defective HW changed, only some dCache pools affected



Update on Atlas TAG DB

- > Database set up (Oracle RAC 4 nodes)
- > Servers ready for upload tests
- > ~3 TB space reserved on NetApp
- > DESY experts in contact with CERN (Florabela) to fill with data

- > Future (looking into my crystal ball)
 - Setup of Web-Frontend (Atlas/Wolfgang, this summer)
 - Is the TAG DB the way Atlas wants to go? (Or ROOT or something else?)
 - Benchmarks: How does backend storage behave? Does it scale? Do we need something else? How expensive? ... ?
 - Cannot tell now :-)



Action item: Batch Monitoring: Waiting time by experiment

- > Code ready and undergoing testing
 - Counting of waiting jobs corrected
- > Time estimate: Before next NUC



Action item: Autoreply for helpdesk tickets

- > No autoreply for naf-helpdesk@desy.de
 - Supporters are fast in the first contact, so no benefit seen
- > Autoreply for naf@desy.de is (conditionally) possible if NUC wishes
 - We propose a reply to the requestor with the same mail as the request, having the RT-number in the subject of the mail
 - Caveat: RT system is undergoing changes in the internal mail handling. One outcome might be to disable autoreply because of potential mail loops.
- > We are currently investigating how users can see their own tickets. Biggest challenge is authentication.



Action item: SL5 sysname

> No standard for sysnames

- CERN e.g. sticks with the default `amd64_linux26` for SLC4.7 and SLC5.3
- So should the NAF

> Experiments should use other techniques

- like e.g. Atlas already does, more details ask Wolfgang

> Alternatives:

- `lsb_release -rs`
- `lsb_release -a`
- `cat /etc/redhat-release`
- `cat /etc/issues`

> Further reading: <http://grand.central.org/pages/numbers/systypes.html>



Current discussion: SL4 plans

- > Latest Kernel (shipped with SL4.8) incompatible with Lustre over native IB
 - Temporary help with 4.7.z-Kernel, no durable solution
 - What to do when newer Kernel is needed? (Security, stability,...)
- > A1: No Lustre mount for SL4 nodes
 - As already indicated in June NUC
- > A2: Local Lustre mount via TCP over Ethernet
 - Much slower, only client reconfiguration needed, server stability unclear
- > A3: Local Lustre mount via IPoverIB
 - ~30% slower, needs server reconfiguration (+downtime), needs testing beforehand
 - Earliest mid of September
- > Additional constraints:
 - Any alternative should not affect the server stability or SL5 running
 - Systems needs to be up-to-date (security-wise)
 - Limited manpower, do not want to spend it on SL4



SL 5 migration

- > DESY-HH CE: Migrated half of the resources to SL5

- Well used by CMS

- > Statement minutes July NUC meeting:

After some discussion the NUC agreed that 90% of the batch resources will be switched to SL5 shortly after the next NUC meeting in August. If contrary to expectations major concern from users show up the migration plan might be altered at the next NUC meeting.

- > Our proposal: Do this during switch downtime (see later)

- Queue default will be SL5
- Interactive WGS default will remain SL4 (can be changed to SL5 for experiments individually when needed)



Longer switch downtime

> Central switch for Grid / dCache / NAF in Hamburg:

- One linecard needs to be changed (spontaneous resets with potential AFS data loss)
- Firmware needs to be upgraded in order to enable logging
- One hour downtime expected
- Affected: All DESY-HH (Grid-WNs , dCache pools&doors), all NAF services

> Other work needed / planned:

- DESY-SE in Hamburg upgrade
- Several pools upgrade in Hamburg (Atlas, CMS, DESY)
- Kernel updates...
- Lustre server update in Zeuthen
- (Migration to SL5?)

> Due to holiday / experiments / external constraints:

- Propose 26.8. (Wednesday) for intervention



Proposal downtime

- > Three days before 26.8. - 26.8.: Stopping respective queues
- > 26.8. 6:00: Start switching off all NAF machines / NAF-AFS / Lustre
- > 26.8. 8:00-9:00 Work on the central switch
- > 26.8. 9:00 DESY-HH CE is available again
- > 26.8. 9:00-12:00 Booting up NAF / AFS / Lustre (including work on Zeuthen Lustre)
- > 26.8. 9:00-16:00 Upgrade of HH-SE for DESY-VOs: Unavailability
- > 26.8. 9:00-12:00 Some dCache pools of DESY-HH Atlas & CMS SE are temporarily unavailable

