

NAF Status

Yves Kemp

NAF Status Report
NUC, 19.11.2014

Issues since last NUC

> Scheduler issues

- Each job writes a script file to one AFS directory
- An AFS directory can contain a maximum of ~64k files
- ... not accepting any more jobs
- Solution: Put the directory to a local storage
- Consequence: Cannot use failover batch master easily

> Currently ongoing:

- Maintenance work on power circuits (“Not-Aus”) in CC4
- ... hope everything goes well

> Have found some systems with erroneous RAID-0 disk configuration

- One by one reinstallation ... unfortunately also some WGS

> gLite UI for SL5 from DESY in AFS is history



Issues since last NUC

- > Questions about #job of different VOs
 - Could be explained by higher memory needs from one VO
 - Some users explicitly asked long.q@sld5 ... but only few resources available
- > Long-running / high-mem script malfunctioning
 - Typo in the script ... fixed
 - BTW: Now responsible naf-VO-support email list is CCed
- > SONAS issues:
 - Sometimes SONAS hangs for some time (some minutes)
 - Was finally tracked down to a bug
 - First half is fixed in the code of minor version introduced Friday last week
 - Second half of the fix only with next major version
- > Login/AFS troubles Monday evening: FibreChannel work with unexpected side effects



Batch migration to SL6/Puppet

- > ... was announced for last week
- > We postpone this upgrade until SL 6.6 is out and some mayhem has settled down
- > ... probably next week of Nov. 23rd
- > Will announce this shortly before we begin
 - ... but decommissioning of ~60 SL5 nodes has started already to make space for new hardware
- > We will introduce a switch `-l puppet=false` for the legacy configuration
- > `-l puppet=true` will be the default
- > We will reinstall SL6 machines as soon as possible, in the meantime, the legacy nodes can be used



dCache news

> End of October: CMS SE

- New hardware for head nodes
- New golden dCache release 2.10.9
- Testing NFS v4.1 access from Grid, currently starting small

> Today: ATLAS SE

- New hardware for head nodes
- New golden dCache release 2.10.9
- No NFS v4.1 tests planned

> DESY SE: No date planned, beginning December possible

- Belle and DPHEP: Some additional space to be commissioned within next days

> ATLAS and CMS: Currently no space issues observed



Infrastructure for CMS CUPS School this week

- > Created ~50 school accounts
- > Created 6 WGS
- > SONAS File Set



Some technicalities about Linux-Groups and NFS

- > You know the UNIX ACLs: User/Group/Others
- > You know that there is a primary group, and people can have secondary groups
- > Sharing data within a subset of people can be done using a special secondary group
 - E.g. af-atltop, ...
- > Caveat: Cannot be extended indefinitely!
 - The NFS client tells the server which secondary groups a user is in
 - The NFS server decides upon this whether access is granted
 - However: The client only transmits a maximum of 16 secondary groups
 - ... if the 17th would be the first match: Bad luck
- > Currently not a problem ... but be warned!
 - Do not plan to create fine-grained ACLs with secondary groups and NFS
 - Alternatives to NFS v3 exist, but are not within reach for us



Some technicalities about shells and environments

- > Problem: Depending on exact job configuration, the command “module” might not be available on Puppet SL6 nodes
- > Why? Several mechanisms to source user and system profile files, not all run in all configurations
- > “module” is a shell function, and needs to be sourced
- > This can/should be done explicitly within the job in the job script file, then it will work in any case:
 - BASH/ZSH: . /etc/profile.d/modules.sh
 - (T)CSH: . /etc/profile.d/modules.csh
 - ... somewhere in the first lines of the script

