

NAF status

[Andreas Haupt](#), [Yves Kemp](#)
NUC virtual meeting 11.7.2012

CVMFS issues

> Background:

- CVMFS pins metadata catalogue files - Unpinning happens only during unmount
- Using automounter – but unmount might not happen due to occupied WNs
- CVMFS throws IO errors when sum pinned files > cache/2

> NAF cache for atlas.cern.ch was 7.5 GB in size – now 10 GB

- No clear recommendation from ATLAS at configuration time, 7.5 GB seemed OK
- Now increased to 10 GB – recommendation in June ATLAS Cloud meeting
- This is the maximum we can allocate to atlas.cern.ch
- Will crash. Already now size of pinned metadata catalogue files is 5103MB

> Bug report opened with CVMFS developers

> Monitoring of error message implemented, cache size/pinned size will follow



SL 5.8 and new hardware

> New CPU hardware in HH

- SandyBridge CPU based
- 16 systems, 12 core/system, 32 GB RAM, 300 GB local disk
- 192 cores / 2.9kHS06 in total
- ATLAS/CMS/ILC/LHCB: 4/2/1/1

> Ready for production – with SL 5.8 (new chipset not supported by SL 5.7)

- NAF currently is SL 5.7
- SL 5.8 is ready for deployment – test tcx168
- Will use maintenance slot on 19.7. for upgrade of WN and IN.



Sonas

- > General Availability Announcement to CMS and ILC users
 - We see some usage
- > Working with ATLAS to refine their configuration
- > Ongoing Work within IT
 - Quota tools (user view, VO view, alarming)
 - Closer integration with internal monitoring and specific monitoring
- > First issue: Automount sometimes (very rarely) does not work
 - Investigating with high priority
- > Current max space allocation (Sum 400 TB):
 - ATLAS: 190 TB
 - CMS: 151 TB
 - ILC: 58 TB



Lustre future in HH

- > Proposal for Lustre future in HH:
 - Schedule elaborated with December conferences and Moriond in mind
- > Users should copy their data from Lustre to Sonas themselves
- > Leave Lustre@HH read-write until 1.11.2012
- > Mount Lustre@HH read-only from 1.11.2012 - 1.2.2013
- > Shutdown and dismantle Lustre@HH after 1.2.2013
- > Lustre storage cannot be reused as Sonas backend. Some reuse as dCache nodes technically possible.



/tmp full / reports of “slow NAF” / multicore

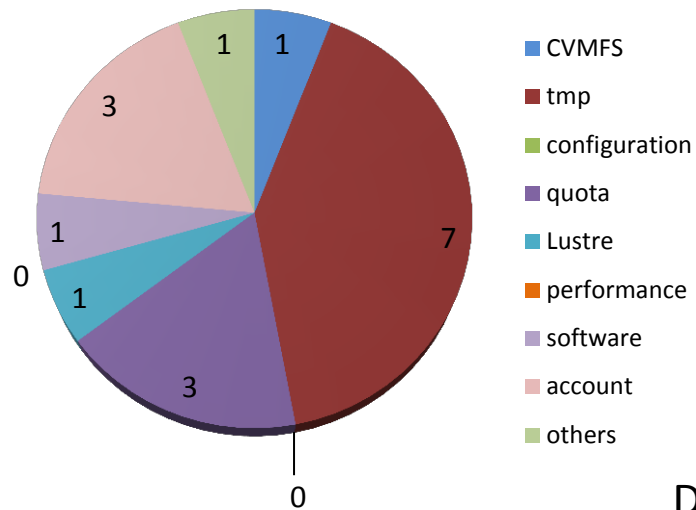
- > One user was writing too large logs – filling /tmp
 - Seen by monitoring
 - Fast manual action (deleting jobs and files)
 - Have to observe how often this occurs to
- > Space problem reported by user of GridControl (CMS) not related to full /tmp. CMS support lists follows issue.
- > One long thread because of “slow NAF”. Error report too vague to follow up.
- > Multicore usage negligible in June.



Naf-helpdesk@desy.de (Jun 1st- Jun 30th 2012)

17 tickets in total

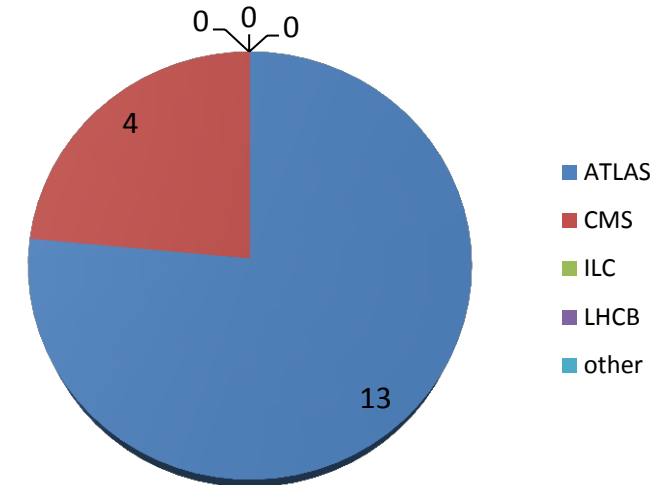
Category of Requests



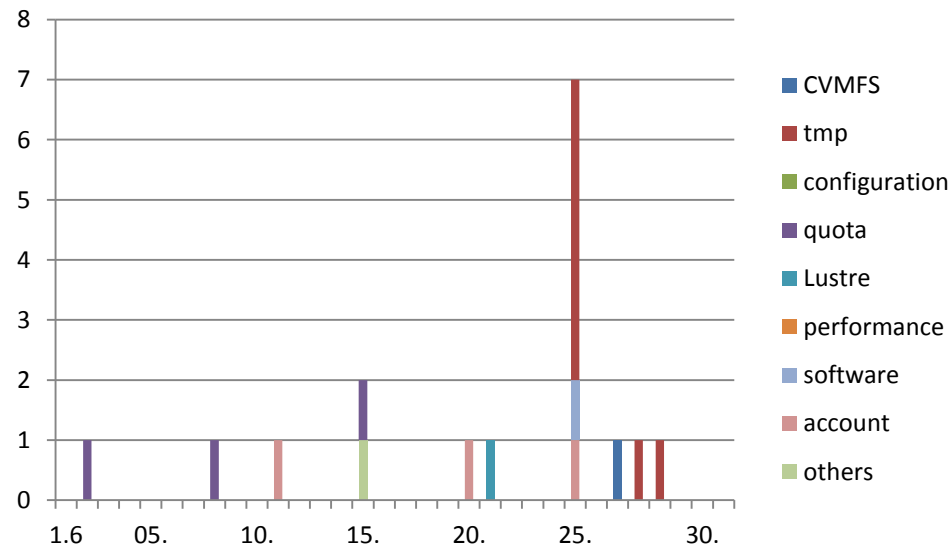
Main Problems:

- CVMFS issue
- /tmp full

Requests by Experiment



Date/Category of Requests



... and something completely different:

- > pub.ifh.de> aklog desy.de
- > pal.desy.de> fs sa \$HOME/private kemp@ifh.de rlidwka
- > pub.ifh.de> touch /afs/desy.de/users/k/kemp/private/testfile

- > Does not work for naf.desy.de as no cross realm trust is configured
 - But it works between /afs/desy.de and /afs/ifh.de
 - ... easier manipulation can be done using *pts* command ... more if needed

