

# NAF Status Report

- Lustre & InfiniBand Status
- AFS Troubles: Status
- Downtime
- Action Items

Andreas Haupt, [Yves Kemp](#)  
NUC 15.9.2010

# Lustre and InfiniBand status

- > Before & during “emergency downtime” on 20.8. in HH:
- > Stripped away all unnecessary IB connections from the switch
  - Identified one blade with constant errors – now retired (HP blade)
- > Reconfigured all HH-tcx-Nodes to run Lustre over TCP/IP
  - No Lustre errors since
  - But: Smaller bandwidth, to be shared with AFS and dCache
  - ... it is clear: This is not a long-term solution, will revert back to IB (see later)
- > Found one server to be in a “funny” state
  - Only explication is a weird order at the initialization – which should not happen
  - ... but we found a good way to check whether LNET (the Lustre network) is OK: Can the machine ping itself over LNET?
  - Did not have this problem since
- > In short:
  - Lustre is working fine since
  - IB seems to be OK (without any load, though...)



# AFS Problems

## > The picture:

- AFS completely hangs for a few minutes, than works again
- Side effects: Login does not work, batch server stutters, registry actions fail, ...
- Does not affect one single AFS file server, but whole AFS instance

## > No correlation found to AFS versions, AFS volumes, users, site, groups, other logs files...

## > First contact with developers: When problem occurs: Set of debug commands

- Check runs once per minute, if problem occurs, debugging takes ~20 minutes
- Hope to get insight from AFS developers, our AFS experts currently meet them at an AFS workshop.

## > Further steps discussed when we know which options we have

## > Just out of curiosity: We see most problems at 7:44, (9:44), 14:44 and 22:44. Do you know of anyone doing something then?



# Downtime

Two parts:

- > Zeuthen: Complete shutdown of whole CC from 16.9 late afternoon – 19.9 ~afternoon
  - Work on (mainly) power, cooling and central router, dCache upgrades...
  - Affects all services (dCache, Grid-CE, ZN-Lustre, ZN-WNs, AFS-Volumes located in ZN)
  - All CMS and ILC AFS-Volumes located in Zeuthen migrated to HH
  - Most ATLAS AFS Volumes migrated to HH (all user volumes)
  - LHCb AFS volumes located mostly in ZN, not migrated (the “big” data could not be migrated anyway)
- > Hamburg: Shutdown 17.9. from 8:00-13:00
  - Upgrade IB firmware on DELL hardware
  - Upgrade to Lustre 1.8.4 on clients and server (bug fix release)
  - Switch back one blade enclosure to InfiniBand (no user access in the beginning)
  - Add new space for ATLAS and CMS Lustre instances. (32 TB each)




# Naf-helddesk statistics (7.8-10.9, Benjamin Kahle)


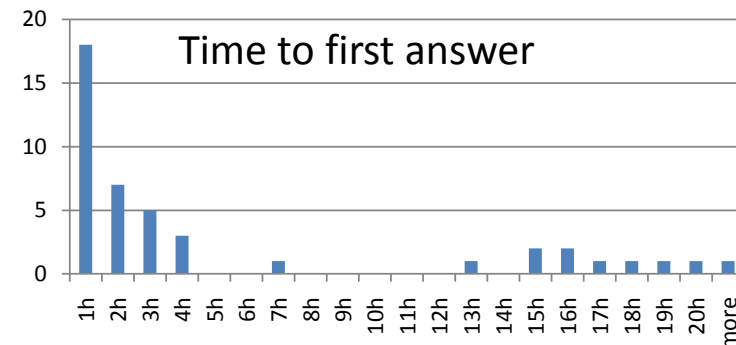
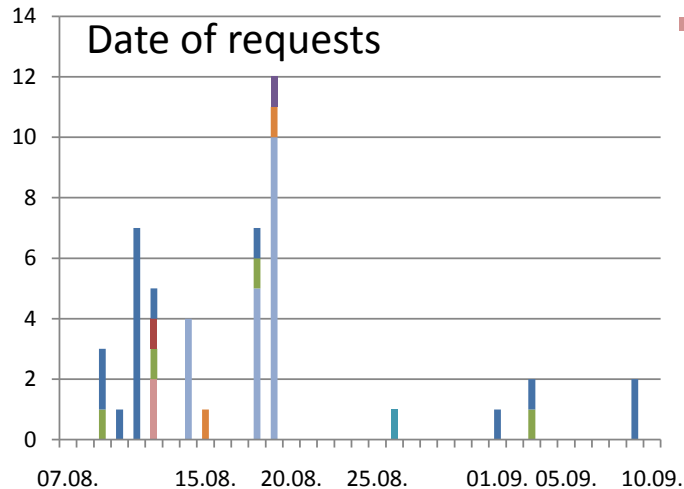

46 Tickets in total

Dominated by AFS  
and Lustre Problems

Category of requests



Requests by Experiment



One wish from helpdesk: If naf-EXP-support is cc'ed, please make sure that haf-helpdesk stays in th the loop



id	who	description	status/comments
old-1	NAF/ATLAS	CMT problem	Wolfgang and Stephan Wiesand look into it, <i>not yet solved</i>
1005-1	NAF	present SL4 work group server usage	show usage at next NUC
1005-2	ATLAS/CMS	plans for SL4 work group servers	talk to users
1005-3	NAF	email notification for /scratch monitoring (instance full)	
1005-5	NAF	automatic Lustre space clean-up	expert meeting foreseen
1006-1	NUC	CPU, dCache, Lustre resource requirements	needed for next PRC
1006-2	NAF	SGE queue for downloads/IO?	<i>still in discussion</i> , use login to WN via qrsh; add to the documentation
1006-5	NUC	ideas for accounting	use Oracle database at HH (wait for Yves)
1006-6	NAF	dCache upgrade time line (10 GE in HH)	is <i>40% done</i> , done until fall
1007-1	NAF	setup Lustre groups for ATLAS, CMS, ILC	done for NAF registry, quota in Lustre, some questions remain
1007-2	NAF	report on help desk tickets	at <i>every NUC</i>
1007-3	ATLAS/CMS/LHCb	ICHEP review; get feedback	

- > Old-1: CMT problem: Wolfgang knows more
- > SL4: see next slide
- > SGE docu contains hint for download/IO operations
- > /scratch notification: When group quotas installed
- > Quotas: tested for CMS: working. Yves will contact ATLAS and ILC
- > Resources & Accounting: New task force initiated by J.Mnich



# SL4 plans

- > Load of 1:
  - Gv running wild, killed this morning
  - Otherwise: tcx030 is bored
- > Logins: List is sent to ATLAS and CMS experts

