



NAF status

... with a focus on network troubles

Andreas Haupt, [Yves Kemp](#)
NUC February 2012 meeting



Network outages 10.1.2012 and after _ 1

Plan:

- > Upgrade core router infrastructure in main computing room
 - Create additional 10 Gbit lines and increase backbone bandwidth
- > SW-Upgrade core router in office network
- > Integration of Grid/NAF/dCache networks in common datacenter infrastructure
 - No separate router, all networks available in the whole datacenter
- > Point 1+3 important for the NAF!
- > More details: CUC presentation by Kars Ohrenberg

<https://indico.desy.de/getFile.py/access?contribId=5&resId=0&materialId=slides&confId=5327>



What went wrong? The NAF view


- > NAF not affected by general outages during 10.1 – in downtime
- > Evening of 10.1: Opened NAF conditionally, but oversaw that WAN problem existed (only tested from internal network)
- > NAF fully functional again only 11.1. evening because of:
 - Switching problems in Grid/NAF/dCache networks – solved by reducing redundancy. Needs additional clarification.
 - Routing problems in Grid/NAF/dCache networks, affected WAN and Zeuthen. Peculiarity of the central switch and its ACL implementation, solved.
 - These problems were fixed 11.1 evening – would have been fixed faster without other problems
- > During the morning of 12.1, last clean-up done (enabling Zeuthen nodes and Lustre)
- > Saturday and Sunday, work in the computing center continued, resulting in an interruption on Saturday 14.1 13:00-14:00 (announced on short notice)
 - NAF spared by most other DESY campus problems



Lustre during 10.1 and after

> Moved most of HH Lustre file server to another computing room on 10.1

- Idea: More stable connectivity to fibre channel SAN
- Made possible by network reconfiguration
- Success: We have not seen problems in the fibre channel SAN since
- Problem: 10 Gbit for these servers was delayed due to general network troubles
- Only yesterday 7.2 switch from 1 Gbit -> 10 Gbit was possible
- Users see a faster Lustre response



Network future

- > Not yet in final configuration ... approaching though
- > In close discussion with the vendor Cisco
 - ... and also discussing extensively details of the outages with Cisco ...
- > As far as we know now, no need for another downtime



25.1.2012

- > We announced network maintenance – fortunately no interruption
- > The same day, one server had hardware troubles and needed to be exchanged on very short notice
 - Host for several important virtual machines, including KDC and Hobbit
- > During Hobbit migration, login was broken: Hobbit load needed during login process - ~1hour
- > To prevent job failures during KDC migration queues were disabled
 - Users submitting with “-w e” see an error
 - ~1 hour
- > MyProxy server (for AutoProxy) also affected



Lustre & Sonas

- > Lustre: ATLAS requested space increase – added ~20 TB
 - No additional fileserver: Bandwidth limitations!
- > CMS has a much more equilibrated read pattern than ATLAS: Reads spread over more fileserver, more throughput
 - Periodic data deletion advisable also for performance reasons
- > Sonas: Currently under heavy testing
 - Performance tests, but also:
 - Functional tests like user management, behavior in case of disaster, monitoring, ...
- > Prior to early-bird-usage in the NAF need to
 - Perform software upgrade to version 1.3.1 – available end of February
 - Integrate old Sonas system – “How” is clear, “when” in discussion with IBM
- > Very tentative schedule: 1.4.2012 for selected early-birds
- > Migration Lustre -> Sonas not discussed yet. How would you like to proceed?



Miscellaneous

- > Users sometimes report AFS client problems
- > We see problems on the server side – but cannot correlate them
- > ... under investigation

- > One user killed *all* ATLAS WGS due to high memory consumption (18.1)
 - We cannot do much – OOM-Killer does not really help in this case
 - Better user education probably the best way

- > AFS Grid-UI: Plan to make default new version glite-UI-3.2.11-1.sl5
 - Objections?

- > HH dCache hardware purchases now fully integrated:
 - ATLAS: 500 Free / 1500 Tot [TB]
 - CMS: 400 Free / 1650 Tot [TB]




Batch usage

- > Usage ~75% -> OK for interactive facility
- > Will still keep an eye on it. Currently, HW purchases difficult though
 - AMD Interlagos has heat problems, no further systems delivered until further notice
 - Intel: Reluctant to purchase now if SandyBridge is expected in larger number in Q2
 - Power/cooling/space issues in HH and ZN: Will be eased/solved during next month(s)

- > Multicore since 1.1.2012

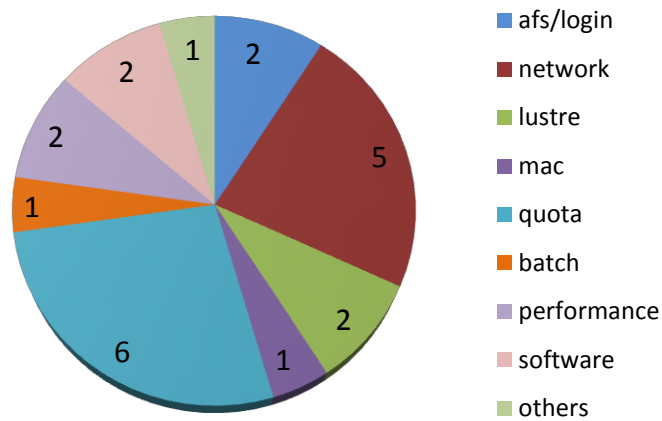
- ATLAS: 289h / 9572 h TOT
- CMS: 1.7h / 45460 h TOT
- ILC: 7h / 5530 h TOT
- LHCb: 0h / 1510 h TOT
- ... we learned recently that CMS is using PROOF-On-Demand – we observe these as single jobs – biasing statistics




Naf-helpdesk@desy.de Jan 1st- Jan31st 2012)

22 tickets in total

Category of Requests



Requests by Experiment



Date/Category of Requests

Severe network problems at DESY

