


NAF Status Report



[Andreas Haupt](#), [Yves Kemp](#)
NUC, 10.7.2013

Problems last month

- > /tmp full – users should write to \$TMPDIR – this is cleaned up automatically after job ending
- > Login troubles and AFS disturbances – see later
- > Java problem with srmcp
 - `export _JAVA_OPTIONS="-Xmx256M"`



Problem reports last month

- > Queue setup NAF 1.0
 - User reporting jobs will not start – transient problem reported by user
- > /tmp problem of one user
 - This time other users reported to misbehaving user directly who killed jobs
- > NAF 2.0 testing
 - Missing RPM
 - Question on Sonas space
- > Slow Sonas reading
 - Not reproducible, no indication in log files
- > (Mis)use of WGS by one user
 - Too many processes, affecting other users
- > Sonas monitoring for Belle @ NAF2.0
 - ... will come



Lustre in Hamburg

- > Announcement to users
- > “1.8.2013 - at which time we will mount the Lustre systems read-only.”
- > “We will switch off Lustre on 1.10.2013, and start decommissioning of the Lustre server and storage nodes.”



Action items

- > 1210-2 NAF 2.0 offer better monitoring for queue usage (per user, waiting time, run time, etc)
 - Working in general on monitoring for SGE batch system
 - Nothing ready to show now – but what you have now will be there in future (as a minimum) 1210-3 NAF 2.0 & experiments reconsider queue setup, scheduling, dynamic or configurable queue limits
- > 1302-1 NAF admins prepare documentation for NAF 2.0 accounts and migration
 - ... still on the ToDo list ...
- > 1302-2 NAF admins clearly document what scratch space means and policy for Lustre and SONAS
 - ... still on the ToDo list ...



Primary / Secondary groups and UNIX-ACLs

> NAF 2.0

- A person can have **one** account but participate in more than one experiment (e.g. ATLAS and ILC)
- As NAF 2.0 is “normal” DESY, internal groups need to decide whether information is shared with external users or not (e.g. small, final ntuples)

> Sonas is mounted using NFS v3 – UNIX-ACLs (user,group,others)

- Need one common group to share data
- When writing data: Need to set the right primary group
- This should be set automatically

> Status:

- Developed script that checks and changes group when accessing WGS through ssh/scp
- Currently undergoing expert user test for bash and zsh, development nearly done for (t)csh
- Batch: Changed BIRD configuration: Jobs run with the primary group at submit time – on a per-job basis



Example




Registry: Namespace H1
Primary Group: h1
Secondary groups:
Atlas, af-atlas, flc, **af-ile**


Example
Config. error




Registry: Namespace AF-ATLAS
Primary Group: af-atlas
Secondary groups: /


PAL

 1g: h1
2g: atlas, af-atlas, flc


 1g: af-atlas (should not login)


nafhh-atlas

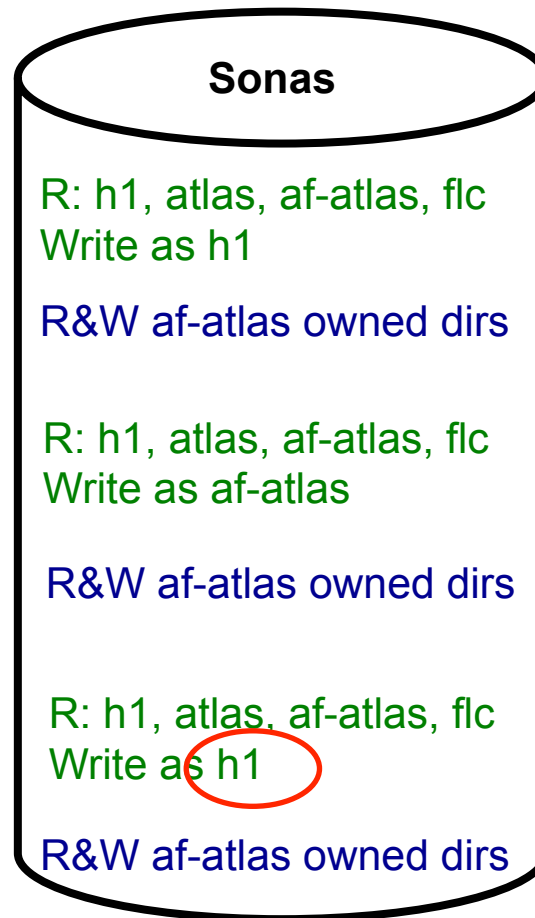
 1g: af-atlas
2g: atlas, flc, h1

 1g: af-atlas

nafhh-ilc

 1g: **h1**
2g: atlas, af-atlas, flc

 Cannot login



Consequences

- NAF 2.0: Each experiment needs own, dedicated WGS
 - On nafhh-x1/x2 (NAF remote desktop) no GID switch – no Sonas mount
- NAF 2.0 users from DESY internal groups need the adequate secondary group – which becomes primary group on the WGS
 - Registration automatisms planned
 - Currently manual scripts to check and straighten things manually exist



> Example for ATLAS:

- `fs setacl -dir /afs/desy.de/group/exp/int -acl usg:atlas rlidwka`
- `fs setacl -dir /afs/desy.de/group/exp/all -acl usg:af-atlas rlidwka usg:atlas rlidwka`

> As people will use this also from non-NAF2.0 machines, an “ls -l” will show “wrong” group ownerships, but they are irrelevant for AFS

Lustre and Sonas

> Lustre: Read-Only mount on 1.8.2013

- Hamburg: All Lustre space (ATLAS: 128 TB, CMS: 100 TB, ILC: 22 TB)
- Zeuthen: ATLAS Lustre space (~50 TB)
- 300 TB in total

> Sonas

- 221 TB of 616 TB in use (that means, bytes written to them)
- Allocated (sum of quota max): 647 TB ...

> 221 TB + 300 TB from migration = 521 TB ... which is below 616 TB

> Nevertheless, we need to increase Sonas. Request for price quote under way

Aging hardware of NAF 1.0

- > Most control hardware out of warranty – some already since 2010
- > Currently no redundancy for login machines
- > Machine running AFS DB evacuated because of recurrent problems
- > ... login problems of users, AFS unavailability



Zeuthen maintenance

- > Zeuthen CC unavailable from Friday 2.8. ~midday – 5.8. midday (maybe shorter)
- > UPS maintenance work

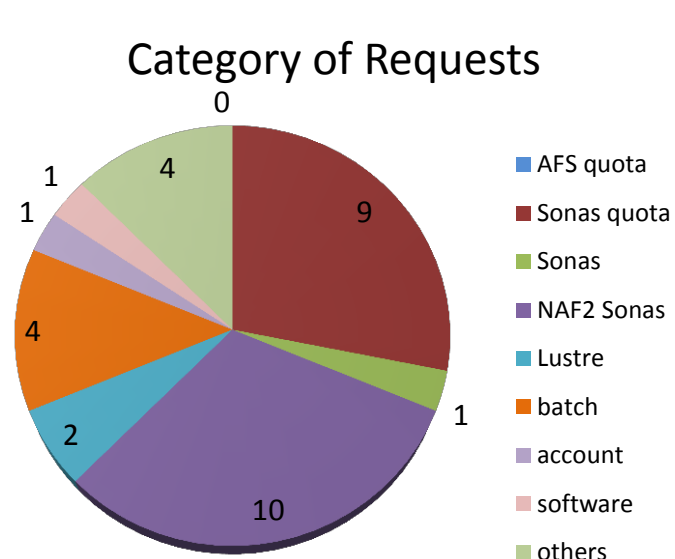


... and the story with the digger hitting the power line



- > Digger hit a cable bridge
- > Maintenance work: power line switched off
- > Computing center 1&2: Only one line
- > This transformer went hot
- > We are on UPS
- > Emergency shutdown to save infrastructure systems
- > According to batteries: ~20 seconds of remaining capacity when power came back...

Naf-helpdesk@desy.de (Jun 1st- Jun 30th 2013)



32 tickets in total

2/3 about Sonas at NAF1 and NAF2

Requests by Experiment

