

NAF Status

Andreas Haupt, Yves Kemp, DESY IT
NUC 9.10.2013

When / how should users migrate?

- > NAF 2.0 is basically ready
- > You can register now – and start working
- > NAF 1.0 will continue to work for some time – but resources will be reduced – soon!
- > After 1.1.2014 major work should happen on NAF 2.0
- > We want to shut down NAF 1.0 in spring 2014



Migration: AFS

- > Personal space in NAF 1.0 : `/afs/naf.desy.de/user/f/foobar`
 - YOU should migrate the space: e.g. using `gsiscp`
- > Personal space in NAF 2.0 : `/afs/desy.de/user/f/foobar`
 - Quota can be adjusted – starting with 4 GB
- > Additional space: XXL-Volume: : `/afs/desy.de/user/f/foobar/xxl`
 - Quota can be adjusted – starting with 8 GB
 - Lesser quality of service
- > DESY-AFS (NAF 2.0) can be easily accessed from outside:
`kinit foobar@DESY.de` (enter your password)
`aklog -c desy.de`
- > Group space: Admins will take care of this



Migration: dCache


- > dCache migration trivial – there is none!
- > Access to dCache with the normal tools. Software might need to be adapted to slightly different environment on NAF 2.0
- > Data path and access method remain identical



- > dcTools
 - SL5: Maintained for short time (until SL6 migration is done)
 - SL6: Working on mounting dCache read-only using NFS 4.1 - make dcTools obsolete

Migration: Lustre

- > Lustre will NOT be migrated to NAF 2.0
- > Lustre decommissioned since 1.10.2013



- > Only one late-comer so far



Migration: Sonas

> Some theory: Sonas is organized in filesets

- Each person has one fileset per experiment
- Some physics groups also have (usually larger) filesets
- A fileset is a management entity within Sonas



> More theory: User ID and Group ID

- NAF 1.0: “id” uid=1234(foo) gid=1009(cms) groups=1009(cms)
- NAF 2.0: “id” uid=5678(foo) gid=3118 (cms) groups=3118(cms)
- Same person has two UID/GID pairs ... clashes with NFS exports

> Therefore: One fileset can *either* be exported to NAF 1.0 *or* to NAF 2.0

- Users can decide individually when to change the export
- Migration steps by expert: Remove fileset from NAF 1.0 – apply UID/GID changes – attach fileset to NAF 2.0 ... after ~10 minutes, a Sonas fileset is migrated
- Group filesets are technically identical, need more organization among users



Current Sonas status:

- > ILC: Complete migration to NAF 2.0
- > ATLAS & CMS: Only few users migrated to NAF 2.0
 - Some still work in the general-purpose test directories in NAF 2.0. These will disappear at some time!
- > Quota and management status:
 - NAF 1.0 quota management with certificate authorization
 - NAF 2.0 will need Kerberos authorization – under development
 - “Which user has space in NAF 1.0 and NAF 2.0 Sonas?” – “Is on a respective WGS in NAF 1.0 or NAF 2.0”
 - Experiment admins can and should manage quotas of their users
- > Currently total Sonas capacity is ~650 TB
 - 200 TB purchased and will be deployed end of the year



Sonas NFS mount and internal / external users

- > DESY users primary group: cms (as an example)
- > External users primary group: af-cms
- > By default: files in Sonas readable by all CMS people:
 - Solution: DESY are given additional secondary group af-cms
- > Make sure that primary group when on NAF 2.0 is af-cms
 - Need to change this on WGS and batch for internal people
 - Performed using a script running at login time on WGS
 - Batch honors primary group at submission time
- > People still can decide differently: e.g. chown ...

**Script working on ATLAS&ILC
WGS**

Debugging needed for CMS



Migration: Workgroup-Server

- > This is the place where you work interactively
- > You can work in parallel with NAF 1.0 and NAF 2.0
 - If you want and/or need to – but not for too long!
 - Your Sonas space is either in NAF 1.0 or in NAF 2.0 ... here you have to decide
 - No “big-bang” migration – you can adapt your scripts, test the environment and do the full migration once you are ready
- > Basically, NAF 2.0 workgroupserver are there
 - E.g. `ssh nafhh-cms01.desy.de`
 - Similar to normal DESY workgroupserver
 - SL 5 for the moment – some SL 6 being deployed
 - NAF 1.0 will not have SL 6
 - More WGS to come – and also a load balancer



- > NAF 2.0 batch is part of the general purpose DESY BIRD cluster
- > BIRD currently consisting of ~1600 cores@SL5 and ~1000 cores@SL6
- > ~400 new CPU cores to be deployed within soon in NAF context and another ~400 new CPU cores in procurement (-> SL 6)
- > Migration of minimum ~500 CPU cores/month to NAF 2.0 envisaged
 - Currently ~3000 cpu cores
 - Relative fairshare for CMS similar as in NAF 1.0 after full migration
- > NAF 1.0: SGE – NAF 2.0: SonOfGridEngine – quite similar in handling
 - Submission scripts should work with little adaption
 - PROOF on Demand should work on BIRD / NAF 2.0

WN & WGS migration plan

- > October: 500 cores -> SL6@NAF 2.0 (2500 remaining in NAF 1.0)
- > November: 1000 cores -> SL6@NAF 2.0 (1500 remaining in NAF 1.0)
- > December: 1000 cores -> SL6@NAF 2.0 (500 remaining in NAF 1.0)
- > January: 350 cores -> SL6@NAF 2.0 (150 remaining in NAF 1.0)
- > March: 150 cores -> SL6@NAF 2.0 (no batch @ NAF 1.0)
- > April: Shutdown of WGS in NAF 1.0

- > Zeuthen currently has ~120 cores in NAF 1.0 batch
 - Moved into “The Zeuthen Farm” or decommissioning (6y old)



What needs to be done?

> Monitoring of all components

- Some already exists – will need to be made available to users
- Batch system monitoring and accounting to come

> More documentation

Current documentation and link to registry:

http://it.desy.de/services/computing_infrastructure/national_analysis_facility___naf/index_eng.html

> Fine-Tuning of WGS and Batch configuration

- E.g. RPMs – usually deployed soon after the request

> ... all this should not prevent users from using NAF 2.0 already now for production!

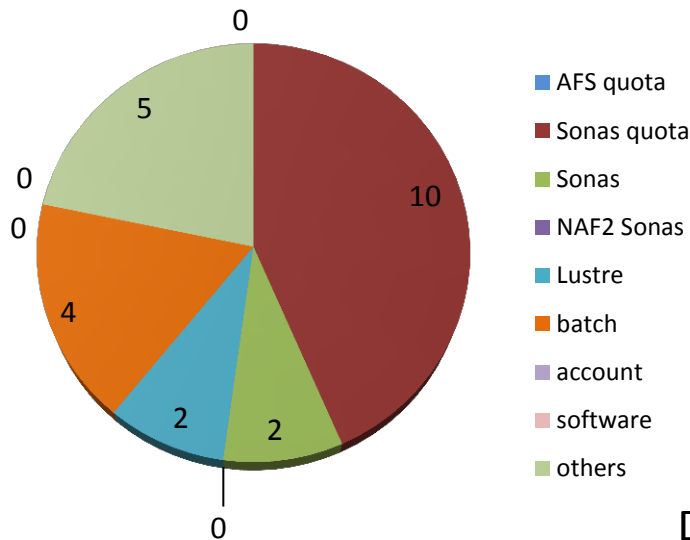


Naf-helpdesk@desy.de (Sep 1st- Sep 30th 2013)


23 tickets in total

Again many Sonas requests

Category of Requests



Requests by Experiment



Date/Category of Requests

