NUC

29.04.2025

Kemp Yves, Beyer Christoph, Dietrich Stefan, Voss Christian
DESY HH

NAF special incidents since last NUC

9th KW (24 – 28 Feb)

- Severe instability of the DUST NFS servers with impact on many NAF components

  - Frequent crashes of NFS servers triggered 90s grace period
  → Allow NFS clients to reclaim locks, no I/O allowed for 90s (ls etc. hanging)

- Primarily affected batch jobs and interactive work on DUST,
Jupyter notebooks affected due to operational DUST NFS dependency

  - no impact on schedulers due to native GPFS mounts

- NAF user experience pretty close to being “unusable”

- Root cause: Known & already fixed issue in latest build

  - ...latest build not yet part of official release (at that point in time)
  → received and applied new build on 2025-02-28

- Apr 16th some jobs went into hold because a small configuration change in the evening was malformed and missed a ‘)’ it was fixed in the morning and the jobs were released

NAF special incidences since last NUC

Fix/workaround for AFS $HOME in apptainer

- since a while we get reports for an apptainer/afs specific problem that occurs and vanishes (which makes it hard to debug for us)

- The symptomatic is that running a container in apptainer fails with an error: We were able to debug the reason behind it to the fact that apptainer does sometimes, spontaneously not like the symlinks inside the path to your home directory in AFS.

- The problem occured before and was fixed at the time (in EL7 context) and we are in contact with auristor, hopefully the outcome will be another fix at one point.

- For now we have a bugfix for it, tell apptainer to ignore your $HOME und explicitly just mount '/afs' by using

- ***"-B /afs --no-home"***

- in the apptainer call. The result will be a fully normal behaviour e.g.

- *[naf-cms14]~% apptainer exec -B /afs --no-home /cvmfs/unpacked.cern.ch/registry.hub.docker.com/cmssw/el7:x86_64 sh -c /bin/bash*

- *Singularity> pwd*

- */afs/desy.de/user/c/chbeyer*

NAF configuration/setup changes

Size of workgroupserver & condor_q

- Running larger computing tasks on the WGS is an ongoing problem causing grief and trouble tickets
- Some of the DASK master processes take 30GB RAM alone on the WGS to host the results
- Decided to give the WGS a makeover and double the resources of the VMs in general
- Was 8 cores 30 GB --> is now 16 cores 60GB (the limit we can offer on the current hypervisors)
  - Further problems will need other measures … as well as situation enables us to do debugging
- ‘condor_q’ cmd is surprisingly costly and one of the most common reasons of job start delays (sched to busy) as people use it in combination with ‘watch’
- Introduced rapper:
  - 10 s delay between 2 calls

```
[chbeyer@naf-cms14]~% which -a condor_q
condor_q=/usr/local/bin/condor_q
/usr/local/bin/condor_q
/usr/bin/condor_q
```

NAF configuration/setup changes

New experimental Jupyter Notebooks

- Especially the BELLE WGS are suffering from people running ‘large sized’ (mainly memory-wise) notebooks

- We are experimenting with individual sized notebooks running on non-reserved resources in the pool

- Test setup on naf-jhub02.desy.de

- Will activate it on the production jhub once last obstacles removed

- If tests with users successful would complete our notebook portfolio

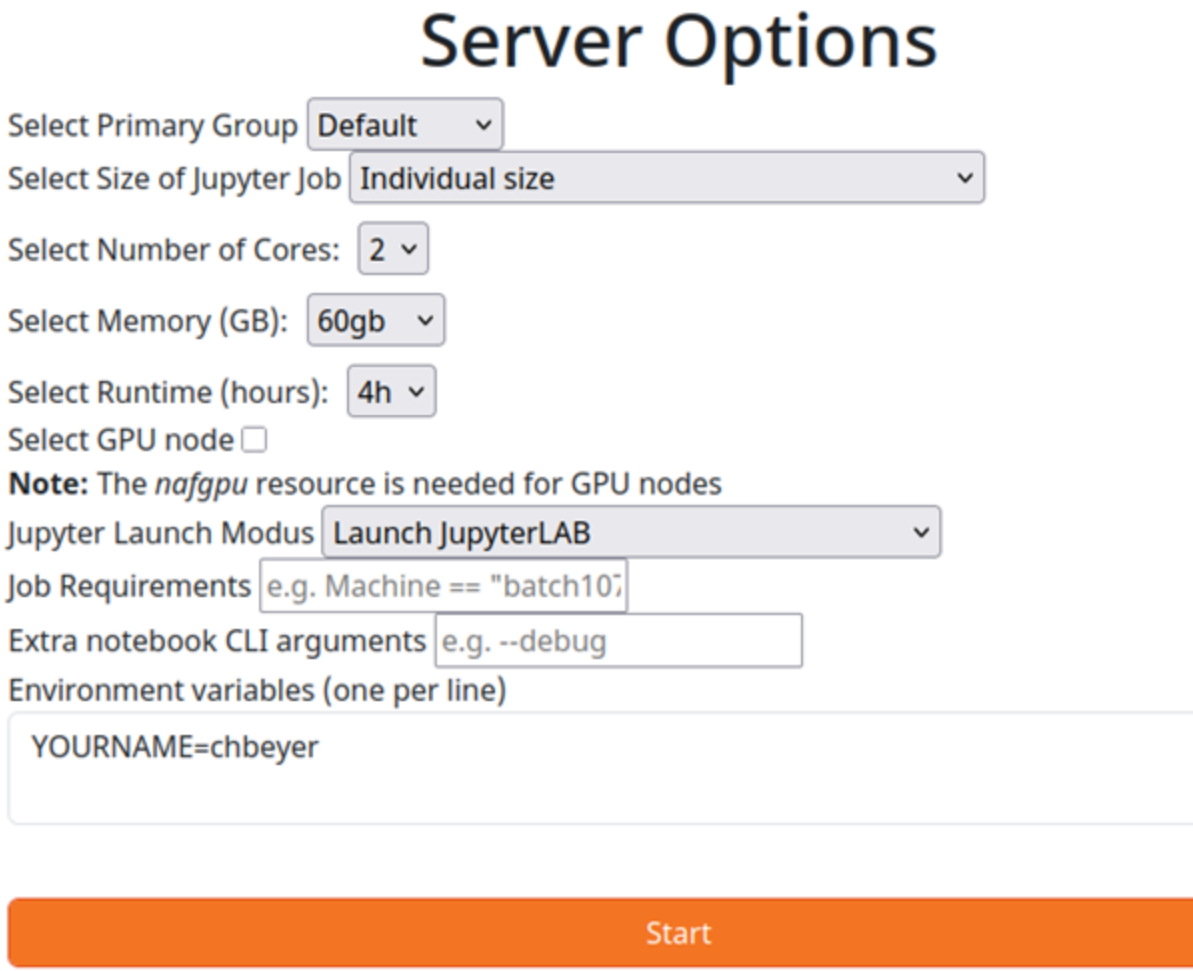

NAF social engineering

“We need to talk” ;)

- Communication leaves room for improvement in both directions
- The Q&A session after the sustainable computing workshop based on the CMS complains seemed fruitful to us (?)
- We would like to offer a regular ‘open office’ type of meeting
  - Roughly one hour
  - Starting with a very short introduction or talk about a topic we think might be interesting
  - Rest of the time open for users questions, problems, ideas
- Some compute expertise from experiment side needed in the room too !
- First tentative date is Tue 13-05 14:00 o’clock

DUST Status

Migration Status & AOB

Migration Status

•Migration finished on 2025-03-05
→ all users and groups migrated to /data/dust structure

•Overall, very few tickets

–One software issue, recommended to rebuild due to complex setup

–Handful “where’s my data?” and “can’t access data of my (former) colleague”

•Removal of ATLAS & CMS symlink compat structure on 2025-05-12

–No symlinks created so far

AOB

•Some more frequent requests to access data from expired accounts on DUST

•Reminder: Lifecycle policy
→ data hidden on account *expiry*

•Proposal to change to lifecycle policy

–Removing dust_user RGY resource has no effect, data remains as it is on DUST
→ dead data on DUST

–Planned: Apply same lifecycle policy, when RGY *resource removed*

–Comments or veto to this proposal?

dCache Status

Upgrade for ATLAS, CMS, Belle II and smaller experiments – Timeline & Procedure

- When opening a ticket with NAF admins, feel free to add a list of files you're reading to allow us to collect data easier per user/job to increase the entropy and reduce the probability of bottle necks

- All dCache instance in the NAF have been migrated to new hardware and golden release
  - Decommissioning of SRM for CMS, ATLAS, and DESY done

- Migration of the university Tier-2 centres can have effects on NAF with another drain on available I/O cycles
- In discussion to mitigate effects with the universities, HPC sites and dCache developers
- Begin discussions on archival as well as FAIR data setup for ALPS II

Miscellaneous

PRC99

- The PRC report, based on the combined NUC/PRC preparation meeting, has not yet been finalized, some preliminary elemtns of discussion
- Analysis Facilities effort, and DESY involvement therein
- User-IT interaction

Miscellaneous

AI based trouble ticket analysis approach

- It has been reported often that in case of troubles, users do not know whether other users have seen or reported these troubles. We have implemented a *pilot* system for summarzing these:
- The NAF ticket queue is checked regularly for active tickets every 5 minutes
- DESY IT is setting up a pilot system of an on-premises AI assistant (based on Mistral). This AI assistant has access (amongst others) to the NAF documentation.
- The initial request of the user in the NAF support queue is given to the AI assistant, with the following task: “Summarize the content of the issue, remove all personal information (e.g. user names), and propose a solution according to the documentation”
- The list of active tickets thus summarized is pushed to the NAF documentation page
- https://docs.desy.de/naf/tickets/