

# NAF Status Report

Untertitel der Präsentation



Yves Kemp  
NUC, 04.09.2019

# Problems with storage last month

- dCache (well, problems related to dCache, not dCache problems per se)
  - Known problems with SL6 NFS clients, leads to overload of the servers
    - Seems to be the major issue with current ATLAS dCache troubles
  - EL7 clients act better, therefore migration to EL7 should be a priority
- DUST
  - Once Ganesha (=DUST NFS Server) down. Restart. OK
  - Once heavy load (2 Gbyte/s = max) on the systems serving the WGS

# Getting rid of SL6

- Container provided by IT:
  - Turns out to be more work than expected, still some background things to clear.
  - Not available in the near future. On the agenda for further future, but will not solve all our SL6 problem now
  - Would not solve the WGS problem anyway
  - Might even delay a real solution, since the pressure for a real migration is away
  - User-started singularity will work of course
- Proposal:
  - Migrate more aggressively batch cores to EL7. E.g. ~1000 cores in the next 1-2 weeks



# Getting ready for EL7

- On the WGS: Proposal
  - Installing half of the WGS with EL7, keep half with SL6 (or 4:2 or other ratio)
  - EL7 systems would be available under alias e.g. naf-atlas.desy.de
  - SL6 systems would be available under alias e.g. naf-atlas-sl6.desy.de
  - Submitting to HTCondor would be
    - SL6 WNs as default when submitting from naf-atlas-sl6
    - EL7 WNs as default when submitting from naf-atlas (=EL7)

# DUST lifecycle policy

- Active, in the sense that the 180 day „hiding period“ has started
- So far 285 filesets are unlinked, only 2 filesets were asked to be linked again
- Experiment admins can see the status, via [amfora.desy.de](http://amfora.desy.de), e.g.:

	user.af-cms.afriedel	1.0000	0.4612	200,195	user	AF-CMS	23-MAY-17	UNLINKED	26-NOV-19
	user.af-cms.aggleton	6.0000	5.3744	1,037,799	user	AF-CMS	18-APR-20	LINKED	-

# Status of the Scheduler itself

- Stable since last NUC
- Question: Limit #jobs per scheduler (reminder: one dedicated scheduler per (large) VO)
- Answer:
  - We see, that schedulers scale very well for „well behaving“ jobs, up to 100k jobs and more
  - Problems are induced by large numbers of „misbehaving“ jobs
  - In theory, we could limit the number, though not tested, and e.g. error message unclear
  - Maybe better alternative would be “late materialization“ ... Testing with users?
  - Purchase of additional scheduler possible, but costly, scaling bad, ...

# Status Jupyter

- One slot per EL7 node reserved for Jupyter (~70, to be increased)
  - So far physical core, will be changed to additional overcommitted core
  - Also one slot per GPU machine ... Needs GPU resource
- <https://naf-jhub.desy.de/>
  - Chose group. E.g.
- Spawning takes ~20-60s (timeout 2 minutes)
- Available resources:
  - All filesystems
  - 2GB RAM, one core, 22h runtime -> needs experience and feedback
- First beta users (so far): Beamline for Schools
- ToDo: Internal enhancements, condor submission, faster startup

### Jupyter on NAF Options

Select Primary Group: ☒ Default  
Select GPU node: ☐ CMS  
Atlas  
Belle 2  
AF Belle 2

Extra notebook CLI arguments (e.g. --debug)

Environment variables (one per line)  
YOURNAME=kemp

Spawn

Upload New

Notebook:  
Python 3

Other:  
Text File  
Folder  
Terminal