

NUC

19.08.2021

Kemp, Yves
HH

Naf Ressource Status

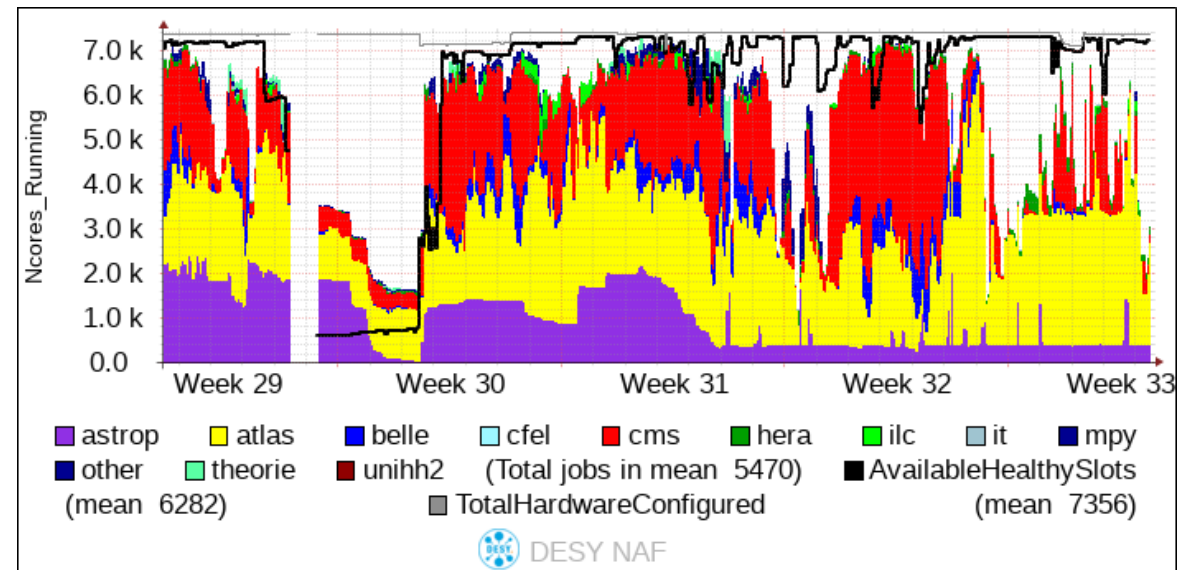
Storage DUST

- Storage nodes are installed
- NFS doors are still in procurement
- Parts of the 'old' hardware shows signs of aging already

NAF special incidences in the past weeks

Urgent Kernel upgrade

- Kernel security issues in week 29 made an aggressive update of all poolnodes necessary
- Jobs with > 2 days running time were widespread over the pool and made fast boots impossible
- Policy-wise we do accept jobs up to 14 days runtime
- Guaranteed runtime is 2 days
- First time we had to enforce this and canceled jobs running > 2 days
- Complete Pool reboot time ~ 4 days



NAF Software

Condor

- major update of condor-version needed due to security issues
 - preparing to upgrade from 8.9.3 to 9.x, currently 2 nodes are running 9.0.4 and behave like expected.
 - will upgrade all WN's first, should be a non interrupting procedure, no reboot, just pkg replacement.
 - Scheduler and negotiator upgrade more tricky and needs to be tested first.
- disk management
 - management in place but it's limits were not enforced as it did not seem necessary until lately someone wrote ~300gb user log & errorfiles due to data sets not being available
 - lots of workernodes ran into disk limitations.
 - Now limits enforced and default set to 20gb
 - No trouble since and very few jobs on hold due to new limits.
 - use: RequestDisk = <size in MB>

NAF Software

Condor

- nested directories
 - previously unnoticed circuit breaker in the condor code when it comes to 'nested' directories currently set to '20'
 - (/afs/desy.de/user/k/kemp is '5')
 - someone ran into it and it took us a while to figure out what happened
 - We could alter it to any other reasonable number now but did not test it yet (opinions ?)
- remove and hold messages
 - working on messages other than <something-in-this-gibberish-condition-evaluated-to-true> if jobs get removed or on hold
 - Works most of the time but still some conditions seem to escape as variables go 'undefined' right after the condition was met.

News from the GRID

Condor

- migrated to Condor 9 on the Batch Nodes and Schedulers
 - Security issue fixed in 9
 - had to migrate CondorCE to 5 (parallel to 9) as well due to dependency issue and pick up a functionality fix
- prepared Condor event logging
 - parsing Condor logs into individual job/daemon/... status update chunks in JSON
 - working on distributing them via Kafka with Elastic Search as initial target for a consolidated cluster view (if it works)
- Kafka consumers for NAF users might be interesting (for information/triggers/...)
 - Realization of authorization not clear at the moment though