# NAF Status Report

**Christoph Beyer & Yves Kemp**

**NUC, DESY 21.2.2018**

# DUST issues since last NUC

- Good news first: No reoccurance of the Oct. issues (partial service failure due to deadlocks on DUST NFS server) symptomatic was very similar as described in a SL6 kernel bug that got fixed around the same time.
- Issues with metadata on DUST in Jan.:
  - HW extension of metadata building blocks with additional SSD enclosure between 2018-01-03 to 2018-01-05
  - Metadata temporary on HDDs caused no user noticable performance impact
  - Very noticable performance impact after migrating metadata back to SSDs
  - 2018-01-08 and 2018-01-10: Short term mitigation, migrate metadata back to HDDs
  - Hanging NFS mounts in same week, most likely due to bad metadata performance => NFS server restart
  - Analysis of storage vendor: hidden hardware fault of internal RAID controller (used by GPFS as very fast cache for small IO)
  - RAID controller replaced on 2018-02-13 and metadata migrated back on SSD a day later
  - No performance problems since then

# DUST issues since last NUC  cont.

> DUST problems this weekend:

- 2018-02-19: nafhh-zeus01/02, naf-hone01 & nafhh-belle02 rebooted on short notice
- Hanging NFS mounts for both pNFS and DUST (shared space in NFS client)
- Unclear situation, as NFS errors from pNFS mount have been logged

> Currently in contact with storage vendor to debug IO errors

# HTCondor status

- Suffered from DUST issues
- PNFS deadlock beginning of february (pending IO blocked kernel on dcache-desy40.desy.de)
- CVMFS problems around 12. feb with stale volumes (also in GRID)
- Implementing more tests and checks concerning fsmounts that will go into STARTD cron later and automatically block nodes with Problems
- AFS token handling still work-in-progress
- Issues with paniking tokenhandler fixed by better signal handling (in busy times >10.000 signals/sec)
- Updated cluster with recent version of tokenhandler on monday
- Fix for a race condition tuesday
- Couple of bugs reported back to the HTC-team (nearly all fixed)

# HTCondor status

- Open bugs:
  - KRB tickets hang around on workernodes too long after job finished
  - STARTD waits forever if no token appears at jobstart (fixed in next release)
- Improvements
  - Two running scheduler in the htc-pilot now (third one is desirable)
  - Implemeted testbed for testing new releases of HTCondor
  - 'Directed acyclic Graphs' (DAG) jobs running
  - Documentation in  Confluence extended
  - User mailinglist is lively
  - Extension of job classadd tested (job keeps track of unsuccessful run attempts and puts those nodes in a 'blacklist' for next negotiation, will be default soon)
- First implementation of hierarchical group quotas running
  - Currently two jobclasses:
  - lite: RequestJobRuntime <= 3h && RequestMemory <= 2048 kB
  - Bite: everything else

# HTCondor status

- In testbed = 3 Job classes
  - lite: RequestJobRuntime <= 3 h && RequestMemory <= 2048 kB
  - short: RequestJobRuntime <= 24 h && RequestMemory <= 4098 kB
  - long: RequestJobRuntime <= 7 d && RequestMemory > 4098 kB
- Quota surplus on 'lite' jobs only makes sure that small jobs run on the 'fast lane' and fairshare is the norm

# Next steps

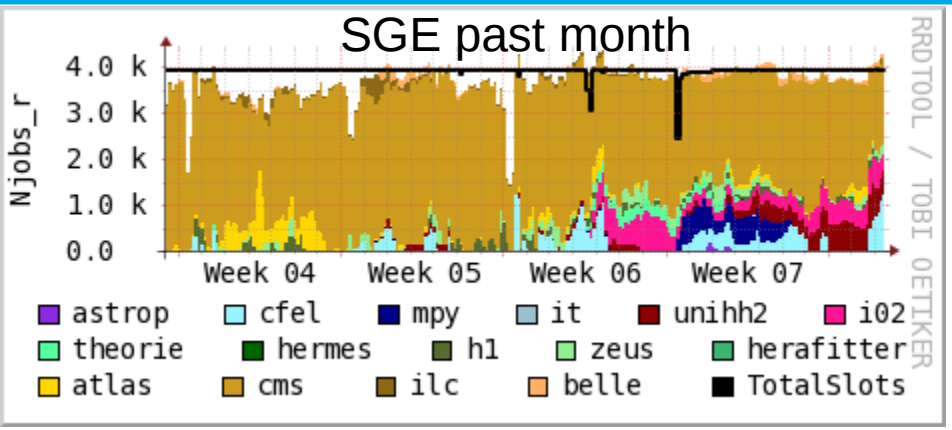


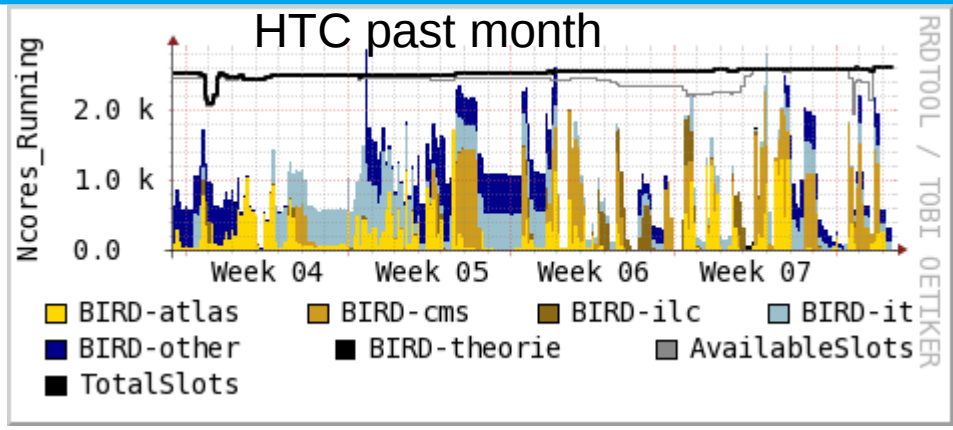


> Further migrate users now? Probably yes.

> How to trigger?

> New nodes awaiting configuration

- Put N% in HTC and (100-N)% in BIRD. What is N? 100?

> Migration plan for the SGE ressources?

# GPUs

- Some systems with GPU purchased by DESY CMS group and UniHH
- Placed in the Maxwell cluster
  - Pros:
  - Enlarging GPU pool, opportunistic usage
  - SLURM scheduling established for GPU usage
  - Software available (easier with EL7)
  - Cons:
  - No DUST available (will never be)
  - AFS: yes, but not $HOME, and K5 token not automatically in batch
  - CVMFS not (yet!) available … but will come
  - No general access to AF-ATLAS
- If need is:
  - We can relocate the GPUs to NAF/BIRD
  - Will stick to EL7
  - Will need to investigate whole-node-scheduling in HTC (we do not foresee another mode of operation for GPU systems)
  - Purchase more GPUs if needed?