NUC

07.07.2022

Christoph Beyer, Yves Kemp, et. al.

NUC, 7.7.2022

NAF special incidences since last NUC

Follow up to NFS overloads

- *Single jobs render workernodes unresponsive for long periods of time*

- New findings
  - Some user request a single-core slot and start multithreaded file transfers (up to 20 threads and more)
  - The workernode runs into load issues and iowait (e.g. load 40.000 iowait 20%)
  - The NFS server gets under pressure as jobs create file requests on a scale that is multiplied by the number of threads e.g. 20 jobs create 400 parallel requests on one volume (hier nochmal mit fileservice kollegen reden ob so richtig)
  - The nodes in question do not start other jobs as load and iowait is too high for reasonable jobstarts
  - As we work with soft limits over time the malicious job gets even more local resources on the workernode as no other jobs start

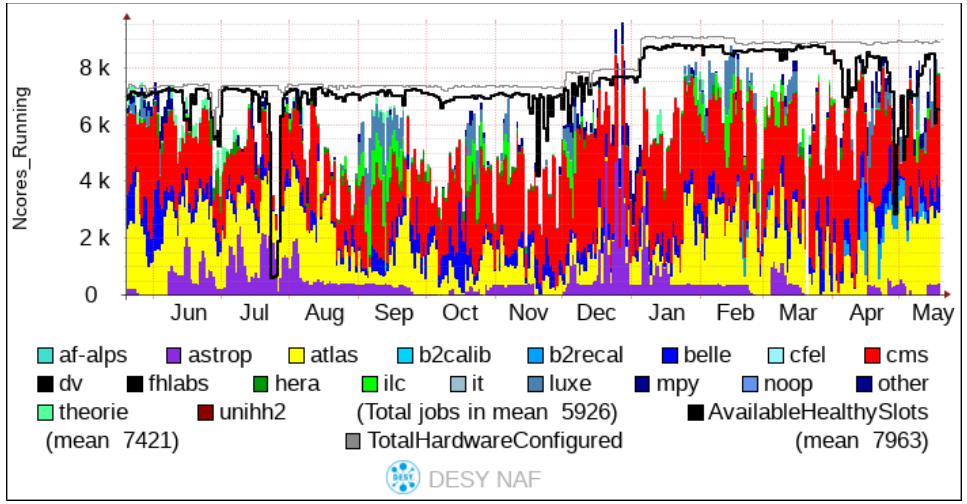

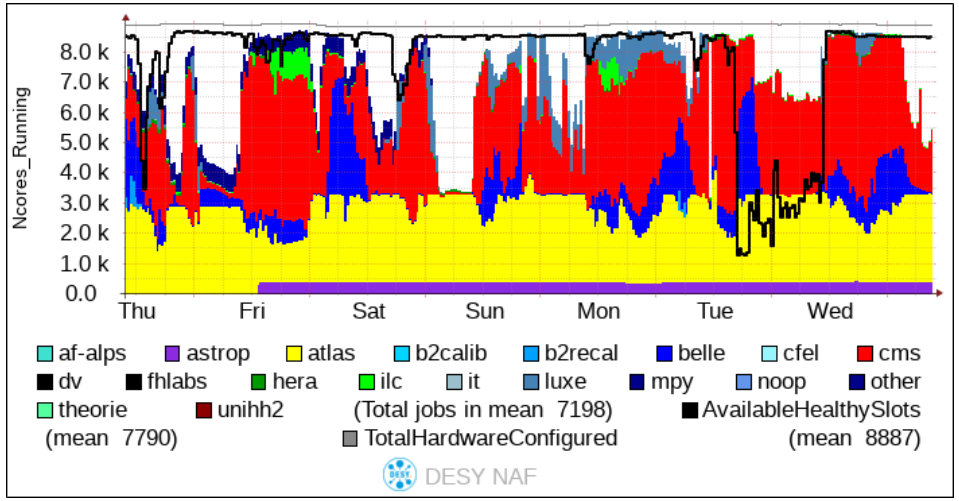

NAF special incidences since last NUC

Possible fixes

- Problem seems python related but not purely
- Belle is a hot candidate for this misbehaviour but again not purely
- Currently very difficult to spot, login on workernode with problems and check processes, then disable all jobs of user
- Traffic is unmanaged from condor view and is not trackable from the condor layer
- Limits per user on nfs volumes come to mind but needs testing what happens when file transfer requests are declined

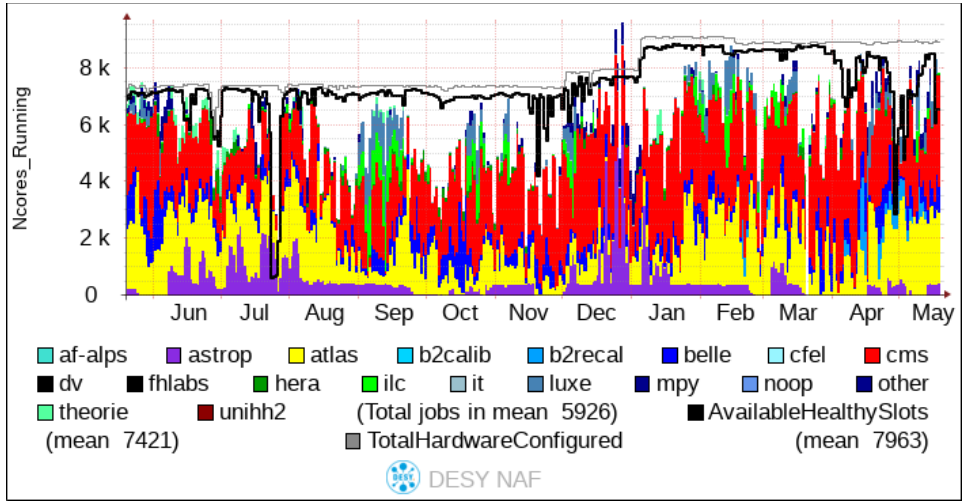

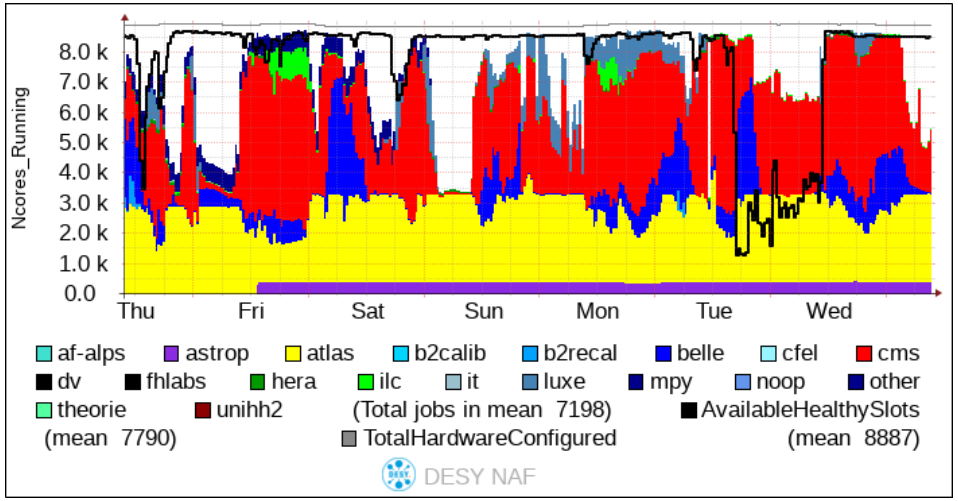

- Symptomatic overlap with problems caused by the fact that IPV6 was not enabled in subnet 76/77
  - Presence of a AAAA record means IPV6 is tried first, fallback to IPV4 after timeout causes different problems for NFS mounts
  - Solved now -> IPV6 enabled

DUST Maintenance 19.7.

- Plan: Upgrade DUST Cluster to newer GPFS version
  - current version no longer supported
- Proposed date: 19.7: needs to be also discussed with Maxwell users
  - Upgrade might take the whole day

- Impact:
  - DUST should remain functional for NAF (and Maxwell) the whole time
  - Performance will be lower during intervention
  - Risk of outage should something go wrong

NAF Storage

GPFS/DCACHE

Updated dCache usage statistics

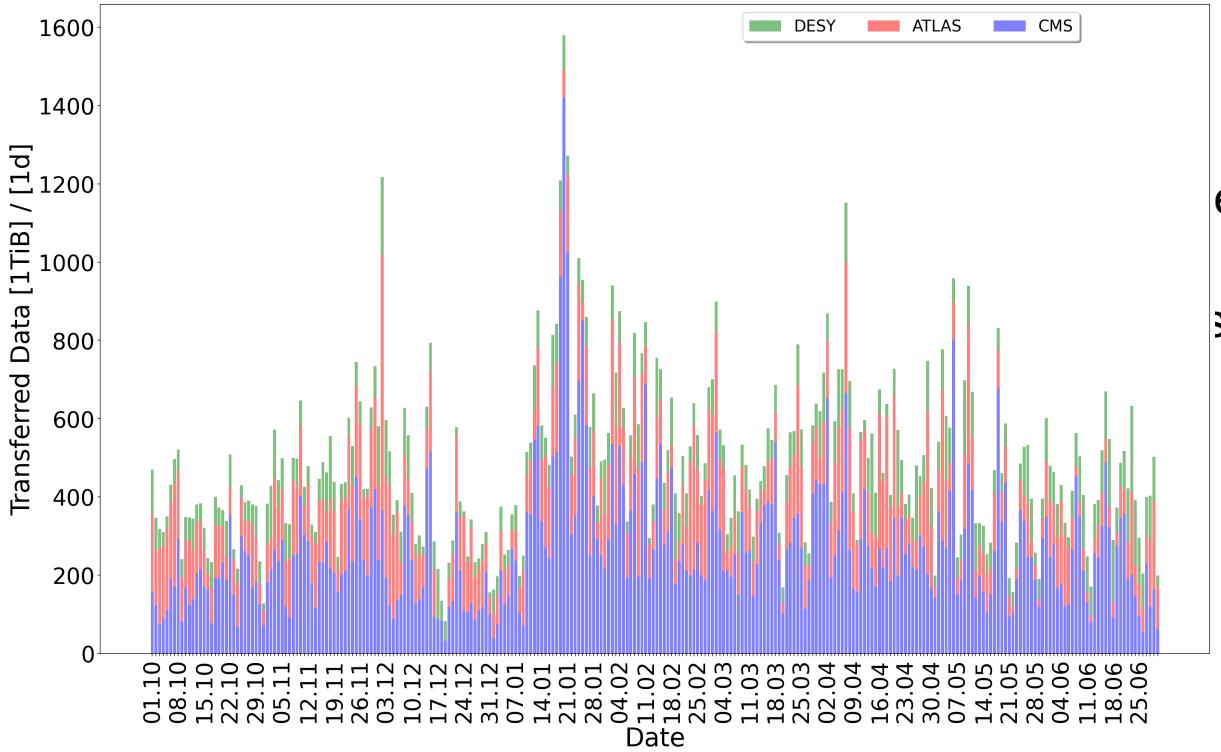

egin

sfers