NAF Status

- SONAS troubles
- Status of SONAS successor
- PNFS problems on WNs
- Status on SGE successor
- AFS status

Yves Kemp
13.7.2015 Hamburg

HELMHOLTZ ASSOCIATION

DESY

SONAS troubles

- Wednesday 6.7. ~13:00
  - „We have a problem with SONAS“
  - It soon turned out to be of larger impact. Call placed with IBM, high priority
  - First announcement to users at ~13:45
- Wednesday afternoon
  - Low level hardware troubles were seen, but unclear error picture.
  - Reboots of affected machines helped for short amount of time only.
  - (Unsuccessfully) migrating data away from one system which appeared heavily affected.
- Thursday
  - IBM experts investigated, and found out that during a previous rebuild (which finished), some blocks could not be read (which was marked as warning).
  - IBM experts found that two blocks were corrupted (we were lucky!)
  - We then were able to pinpoint to two files actually using these blocks (time consuming)
  - We tried to rescue the files, but not possible, needed to delete them to migrate data away

SONAS troubles _ 2

- Friday morning
  - Migration of data from the partially corrupt filesystem was done.
  - At ~10:00, we opened the WGS again.
  - Migration away from the whole fileserver went on until Friday evening. (IO intensive)
  - Over the course of Friday, we slowly brought back full batch capacity again
- Weekend: Stable operation.
- Now / Aftermath
  - Generally, the rebuild operation is routine: ~10% of all Sonas disks exchanged in last 4 years, with only one major problem.
  - Understanding with IBM/SONAS why this error could happen, and why it was only reported as warning
  - Currently some capacity offline (~5 TB) as we took out the affected hardware

SONAS performance plots

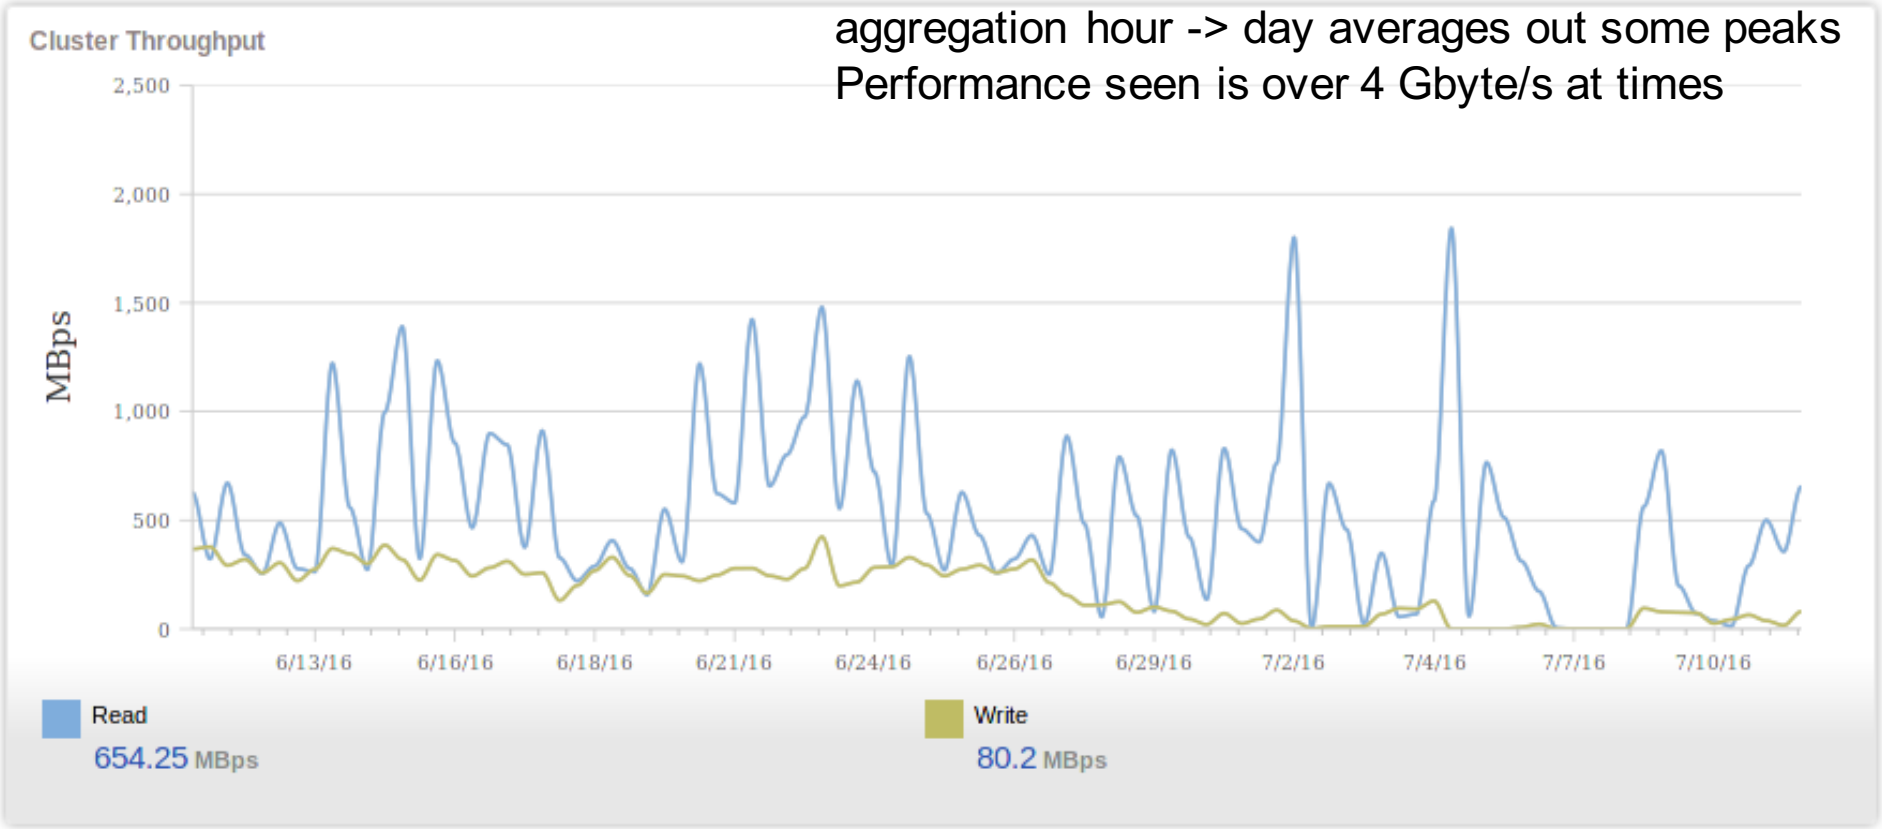

aggregation hour -> day averages out some peaks
Performance seen is over 4 Gbyte/s at times

SONAS successor

- Reminder: IBM GPFS in strategic use at DESY for scientific data
  - Natural to use GPFS filesystem for SONAS successor as *NAF scratch file system*
- What we know and have
  - Hardware from IBM (and some DELL), software GPFS with native raid from IBM
  - ~1.8 PByte raw capacity (~1.3 PByte usable), already delivered
  - Called DUST (well, DUST2 to be exact)
  - Setup needs some time: Currently fighting with dead-on-arrival hardware...
  - Will perform burn-in and initial load tests before pilot production
  - Currently working on quota management migration to AMFORA (the AFS quota mgmt tool)
- Prospected migration plan
  - Start with CMS ... Would also come in handy (hopefully) for the CMS-DAS school in September
  - Migrate the other groups N month later (N>1-2 && N<9)
  - IT does migration for a COMPLETE VO: We copy all data (rsync under the hood). Downtime only for final diff and re-mount. Max. 2 days downtime (e.g. in batch)

... One more thing on SONAS & DUST2

- Policy for data of expired accounts
- We will unlink data once an account is expired.
  - Unlink a fileset: No longer visible, but still existing on disks
- After 180 days: Removal of the fileset
- Better: Users should organize data inheritance ***BEFORE*** account expires
- ... Somewhat in line with what was decided in CUC

PNFS troubles on BIRD WNs

```
Active Internet connections (w/o servers)
Proto Recv-Q Send-Q Local Address
...
tcp        0      0 131.169.84.217:816      141.34.228.53:2049      ESTABLISHED
tcp        0      0 131.169.84.217:868      131.169.64.132:2049     ESTABLISHED
tcp        0      0 131.169.84.217:885      131.169.64.156:2049     ESTABLISHED
...
```

RPC-Port

SONAS

„NFS-Port“

Two dCache pool nodes

Source-IP:Source-Port

Dest-IP:Dest-Port

- A connection to every NFS server that the NFS client speaks to is opened. On a different source-RPC Port
  - One dCache pool node can host many pools – each pool is treated as one server
- The port range is 665-1023 = 358 possible values
  - ~340 dCache poolnodes, with many more pools
- What can happen – and *does* happen
  - Kernel runs out of free RPC ports on the client
- People notice this as „cannot access dCache/Sonas/NetApp on wnXYZ, but on wnABC or on wgsNN“
  - This is **NOT** a SONAS or dCache or NetApp **Server** problem!

PNFS troubles on BIRD WNs _ 2

- Hard to debug ... Intermittent error on WNs
  - The ports are freeed again after some timeout
- Very lucky to have an NFS server developer inhouse! (Tigran)
  - Got down to the root cause of the problem
- Long-Term-Actions: Tigran proposed a modification in the kernel
  - Unlikely that it will make it into RHEL 6 (=SL 6)
  - RHEL 7 might be possible
- Short-Term-Action: Monitoring, Mitigation
  - Monitoring of usage in place
  - Cool-Down script in place: Take a node offline if RPC-port-usage > 80%
  - Not yet done (consequences unclear): Lower the threshold from 665 to e.g. 256

Status HTCondor

- HTCondor is a scheduling and batch system. (Some) main features:
  - Extreme scalability (CERN n*100k cores)
  - Free and open sources, many adaptors in HEP
  - Optimized for high throughput clusters (as opposed to high performance)
- We opt for it as successor of PBS/Torque (Grid) and SGE/SoGE (BIRD)
  - Long-term idea: combine BIRD and Grid farms

Status HTCondor @ GRID

- Currently ~8k cores in pool (using hyper-threading)
  - This is the half of the resources that is under warranty. The other half still with PBS
- Jobs submitted via 2 ARC CE, one condor master host
- Working: Multicore jobs, quotas, installation via puppet, Grid UI and experiment software via CVMFS, monitoring Icinga and grafana
- Todo: More monitoring, housekeeping of WNs, ...

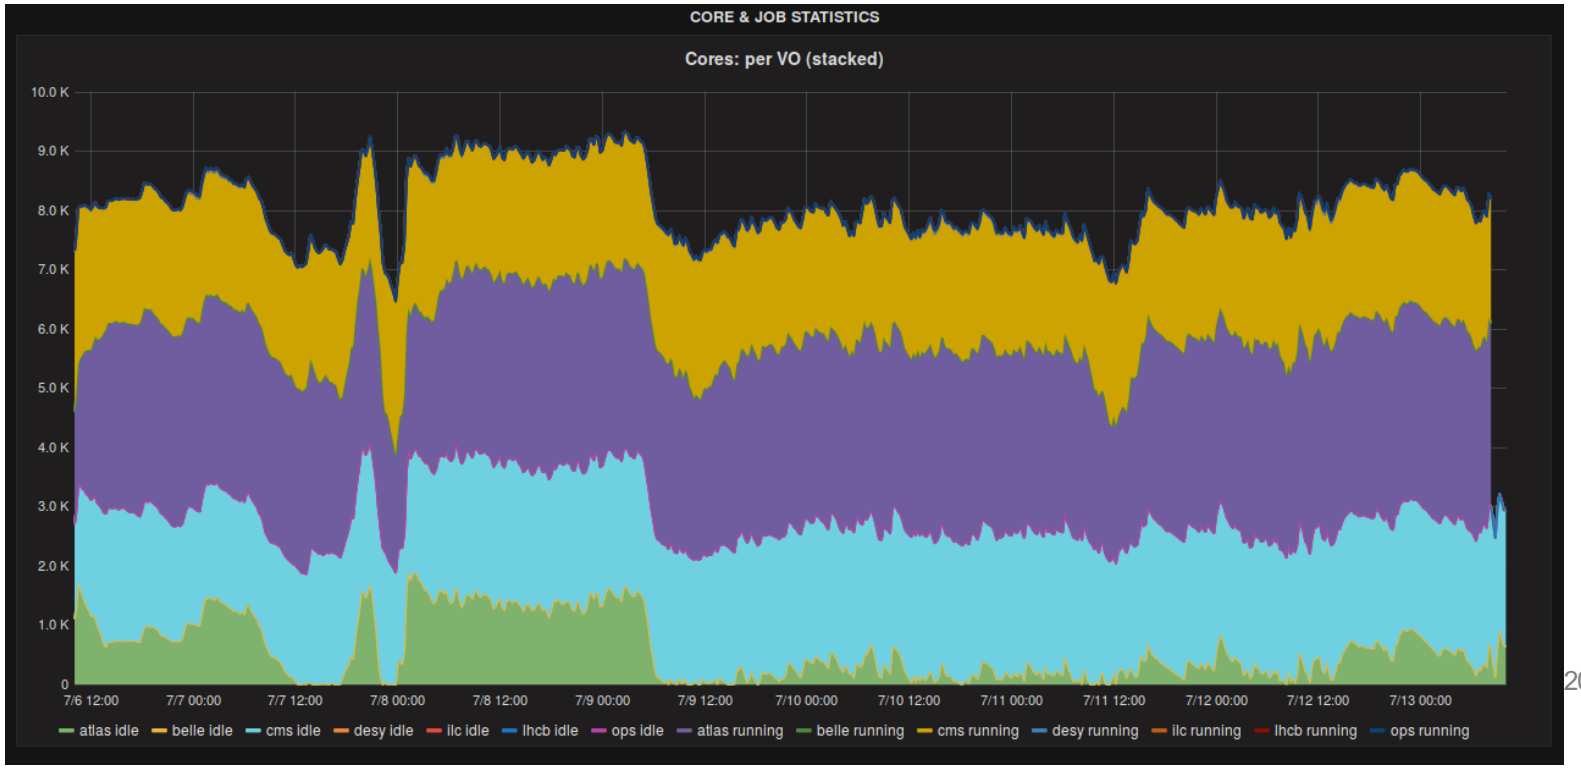

Status HTCondor as BIRD replacement

- BIRD needs Kerberos integration
  - Works in testcluster with a release specially crafted for DESY since ~1 week
  - HTCondor jobs can access AFS (with ACLs). Kerberized NFS possible.
- ToDo:
  - Testing and hardening ot the Kerberos/AFS integration
  - Even more testing and hardening of the Kerberos/AFS integration
  - Understanding fine-grained quotas for BIRD users
  - Substantial test cluster for brave NAF users
- Plans:
  - If CMS-DAS needs batch resources, they will go to HTCondor
  - Subtantial test cluster still this year
  - If things work well, migration end 2016/beginning 2017 in short amount of time
  - SL6 (and EL7) only!

Status of AFS / OpenAFS

- Have you followed the CUC in the last three month?
- In a nutshell:
- OpenAFS project has trouble delivering working kernel modules for Ubuntu
  - There is some danger for RHEL derivatives, but currently we seem to be fine.
- Nevertheless, we must think about what comes after OpenAFS
- A questionaire was sent out to DESY groups. NAF users not explicitly asked.
  - ... But we can/should discuss this also here