

NAF Status Report

Main operational issues last month

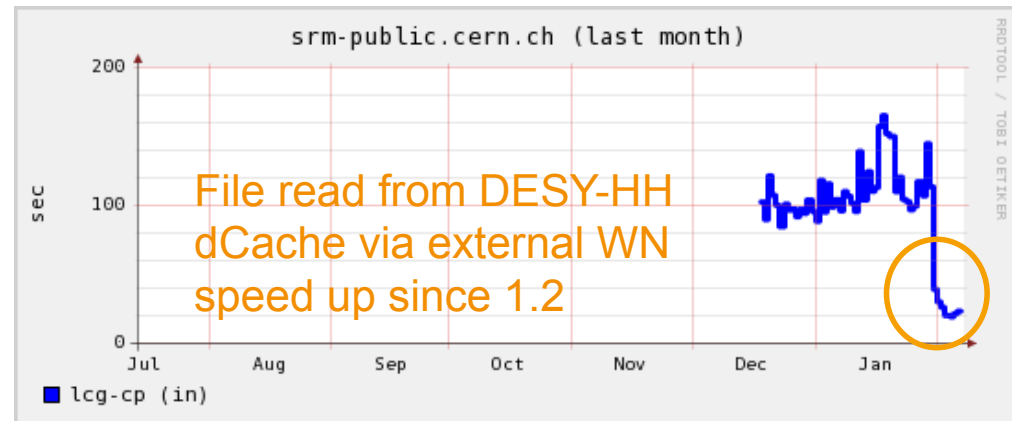
Going through Action Items

Andreas Haupt, [Yves Kemp](#)

NUC 9.2.2011

Network

- > Around 25.1-26.1: 10 Gbit link between HH and ZN down, using 1 Gbit backup line
- > Evening 27.1/Morning 28.1: Failure of one core router in HH, affecting basically all services in HH and partially routing from NAF-HH to NAF-ZN
- > 1.2: Replacement of firewall. From 6:00-8:00 HW installation, some rules had to be adjusted after this time
- > + some minor issues



- > “Future” work: 10 Gbit Infrastructure in HH needs HW-fix (wrong tin-solder used by manufacturer)
 - Will be done Friday [11.2@9:00](#) , takes ~15 minutes, and should be transparent



Batch news

> New batch share after exchange of old HH WN

- ATLAS/CMS/ILC+CALICE/LHCb = 4/2/1/1 (for NAF purchased HW only!)
- Ncores and estimated HepSpec taken into account
- Estimated HepSpec06 Total: 22kHS06

| VO | Old FS (%) | New FS (%) |
|------------|------------|------------|
| ATLAS | 25.6 | 28.1 |
| CMS | 43.6 | 41.6 |
| ILC+CALICE | 4.2 | 5.1 |
| LHCb | 26.5 | 25.2 |

> PROOF & Multicore usage by experiments since 1.1.2011

| VO | Wall-Days | % of all Wall |
|-------|-----------|---------------|
| ATLAS | 794 | 6.7% |
| CMS | 51.6 | 0.2% |

- > Reorganization of AFS fileserver finished
 - Each VO has two dedicated FS (volume replicas might reside on FS of other VOs)
- > Have seen fewer incidents related to AFS and massive job start
 - Is the load better spread?
 - Are first mechanisms working?
 - Are users following advises?
 - Are users avoiding the NAF?

SL 5.6

- > One testing IN: **tcx168**

Two potential problems

- > Tcmalloc was reported to be buggy, LHCb has seen issue
- > [root@tcx168 ~]# `svn --version`
`svn, version 1.6.11 (r934486)`
- > New Kernel already active on all nodes (recognizes 10 Gbit cards)
- > Please test and report!
- > No time scale for deployment yet ...



Testing Job Submission Verification Scripts

> Status

- Current (base-)version: non-destructive
- First feedback: Increase maximum runtime for JSV-scripts
- Only restages IO when JSV is used and IO path is **not** specified (default)

> Next iteration:

- Speed-Up: Rewrite in Perl (instead of bash)
- Next test: **All** SGE-IO path transformed to local path (and later staging)

> Last iteration

- Move from personal AFS-public to local WN location
- JSV used for all jobs by policy



Action Items

| id | who | description | status/comments |
|--------|-----------|--|-----------------|
| old-1 | NAF/ATLAS | CMT problem | No news |
| 1005-3 | NAF | email notification for /scratch monitoring (instance full) | done |
| 1006-6 | NAF | dCache upgrade time line (10 GE in HH) | done |
| 1007-1 | NAF | setup Lustre groups for ATLAS, CMS, ILC | See slide |
| 1007-2 | NAF | report on help desk tickets | See slide |
| 1010-2 | NUC | resource requests | |
| 1010-4 | NAF | statistics on multicore/PROOF usage | See slide |
| 1012-1 | NAF | evaluate different stdout/stderr handling for SGE | See slide JSV |
| 1012-2 | NAF | report on CERN strategy for batch and AFS | Last meeting? |
| 1012-3 | NAF | new fair-share after hardware upgrade | See slide |



Action Items (2)

| id | who | description | status/comments |
|--------|-----------------|---|---|
| 1012-4 | ILC&NAF | develop strategy to manage Lustre space | No news |
| 1101-1 | NAF/experiments | WGS reliability/support/monitoring | See slide |
| 1101-2 | AH/WE/HS | naf_token for ubuntu&MacOSX | No news |
| 1101-3 | NAF | revisit: user information/twitter | See slide |
| 1101-4 | NAF/CMS/ATLAS | move support list to request tracker? | No email received notify YK via mail |
| 1011-5 | ATLAS | written request for NX | |
| 1101-6 | NAF/experiments | do we need a faster x-connection? | |
| 1101-7 | NAF/LHCb | limit number of parallelly running jobs | See slide |
| 1101-8 | experiments | test new job submission verifier | |



Action Items (3)

> Lustre Group quotas:

- New Lustre Client installed, first tests in separate instance successful
- ... but we would need a downtime for testing it ... no downtime at the horizon yet

> WGS reliability/support/monitoring

- Have installed hardware scanner to better detect failures
- Have reminded operators of procedures (boot ATLAS-IN -> Mail to supporters)
- Rearranged IN of all VO, ATLAS++

> Twitter: Our script does not work anymore, Twitter changed its authentication method ☹

> NAF/LHCB: Limit number of concurrently running jobs:

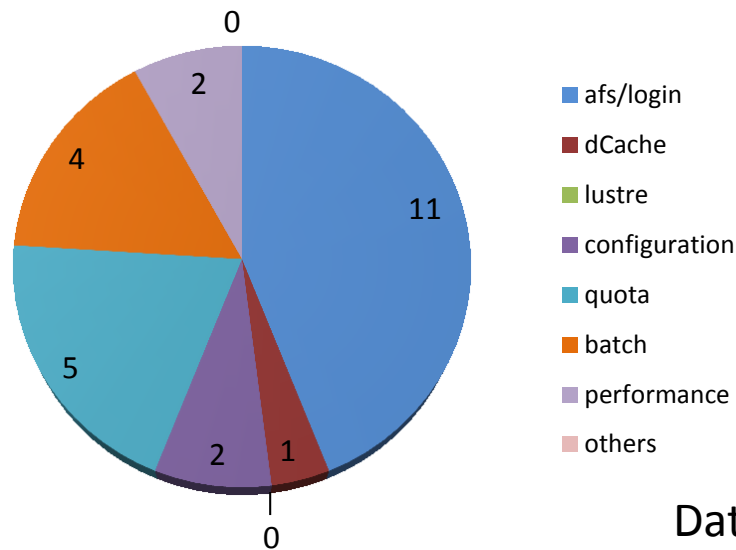
- Needs some final communication between LHCB&Andreas Haupt




Naf-helpdesk@desy.de (Jan. 1th-Jan.31th2011)

25 Tickets in total

Category of requests



Requests by Experiment



Date of requests

