Castor Status

Castor Task Force:

Address with urgency the performance and stability shortcoming of CASTOR impacting the Tier-0 and CAF operation at CERN.

- The team includes : Castor operations , Castor development, members of GD and DES group
  plus close collaboration with the ATLAS operation team
  Team leader : Bernd Panzer-Steindel

- The activity started fully in the first week of April.

- During the last 3 month the focus was to provide a new Castor release with much higher stability. Other Castor activities were treated with lower priority (SRM, xrootd, some part of the support, etc.)

- Every morning a ~30 min coordination meeting, plans for the next 24-48h and adjustment of priorities
  plus several dedicated longer coordination and planning meetings

started with release 2.1.3-3 in April → production release is today 2.1.3-15

- ~40 functionality improvements and changes
- ~35 bug fixes
- improvement of the installation procedures
- improved testing, more test-suites
- move from LSF6 to LSF7
- all stager hardware now on the new NAS Oracle servers

e.g.

- LSF-plugin rewritten
- monitoring rewritten
- load-balancing partly rewritten and externalized
- fixing several problems in the file system space allocation area, including the rejection of jobs when there is no free space (disk1tape0)
- different scheduling of some commands to ease the LSF queue load (PutDone)
- one LSF queue per service-class/disk pool
- fix „broken/residual“ entries in the stager DB

new release deployed,
running the ATLAS T0 exercise
since 8 weeks

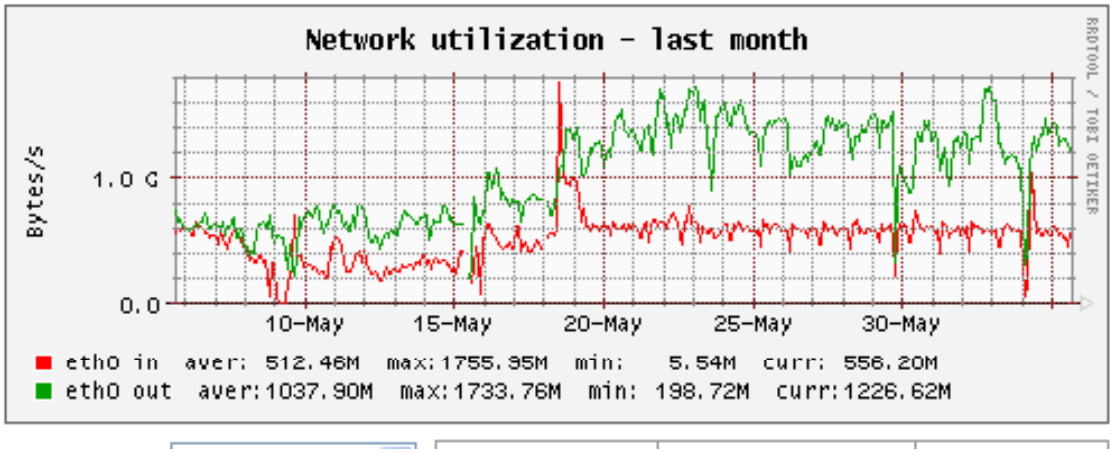

major changes:
- improved the load-balancing
- doubled the hardware
- fixed file system space limit problems
- increased the ATLAS test complexity

50 disk servers, ~300TB
200 worker nodes
16 dedicated tape drives

3 disk pools
(disk1tape1 and disk1tape0)

→ full nominal ATLAS speed for emulated DAQ-T0, reconstruction, AOD merging, tape storage.

→ export on average 500-600MB/s with short peaks at 900 MB/s
still ongoing tests and analysis
the new ATLAS nominal export goal is 1020 MB/s (increased event size) and we can‘t reach this yet (Castor?, sites?)
very complex environment

Test example

deleting of files is a heavy operation, needs some more DB tuning

ATLAS started a major clean-up process
→ 85000 file remove operations per hour

after the clean-up the system absorbed the backlog

system recovered to nominal speed

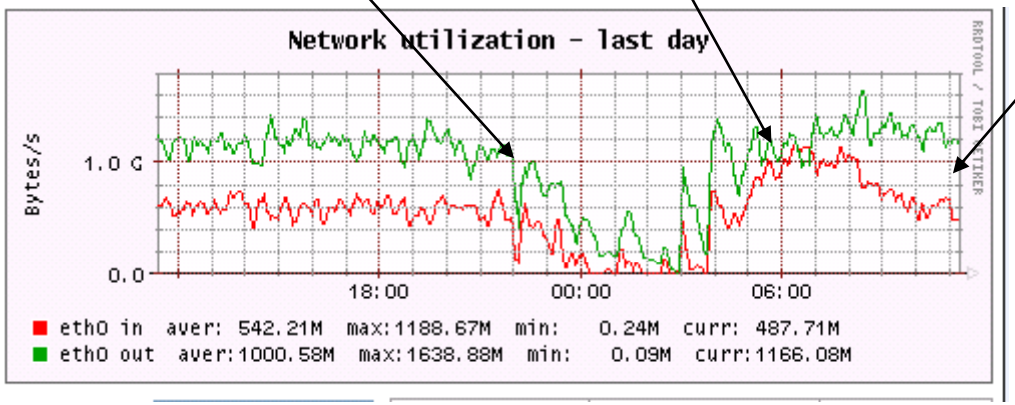

no intervention needed
system reacted well and recovered

Continuous testing of the new release:

- dedicated tests of components
  - → >50000 jobs in the LSF queue
  - → establishing Castor limits, e.g. 15 Jobs/s LSF throughput, 150 requests/s into the DB
  - → optimize the tape writing, try to reach 60 MB/s (average is now ~45 MB/s)
  - → optimize footprints of requests on the DB level
  - → .....

- „interference“ tests

  e.g.
  running the ATLAS T0 exercise
  plus
  8000 read and write streams in another disk pool
  plus
  loading the stager with 90000 query requests per hour

  → no effect on the T0 exercise

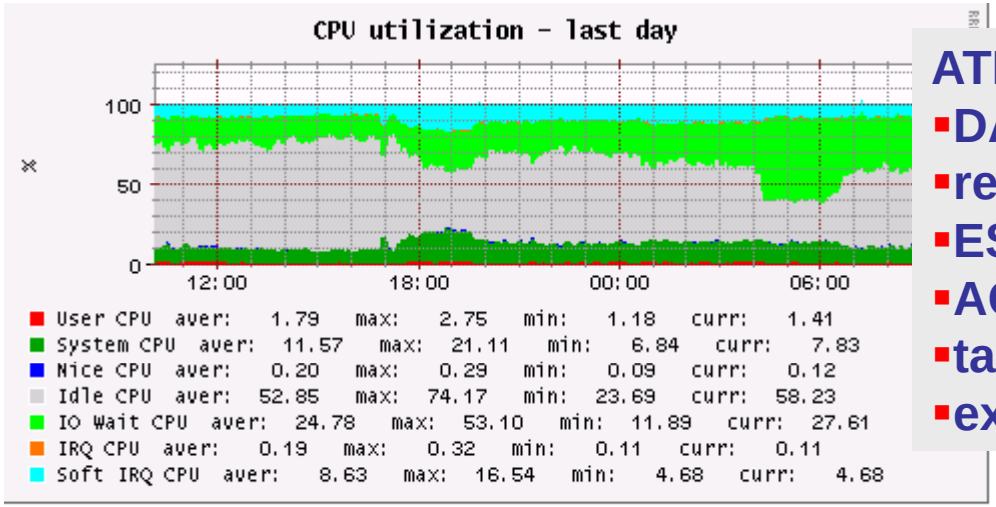

ATLAS T0 exercise running with
- **DAQ emulation**
- **reconstruction,**
- **ESD+AOD production**
- **AOD merging**
- **tape migration for RAW, ESD and AOD**
- **export**

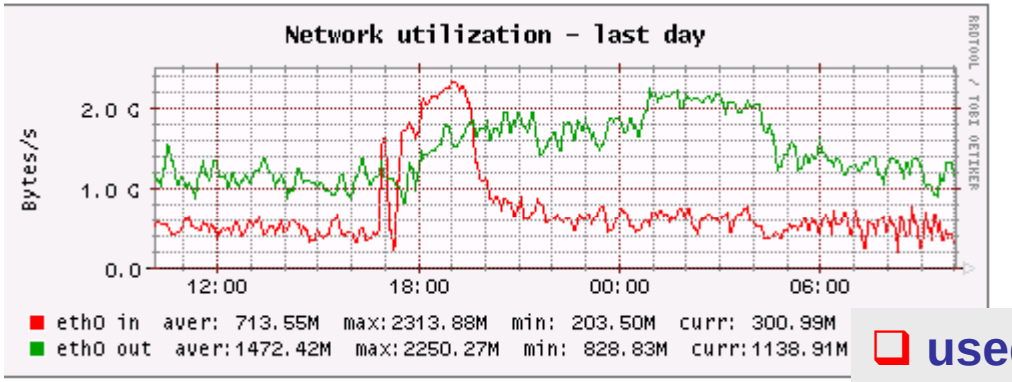

- adding additional streams through Castor showed problems and did not show the expected performance increase
→ Castor load balancing issues

- **used 50 external nodes to generate extra traffic in the t0perm pool**

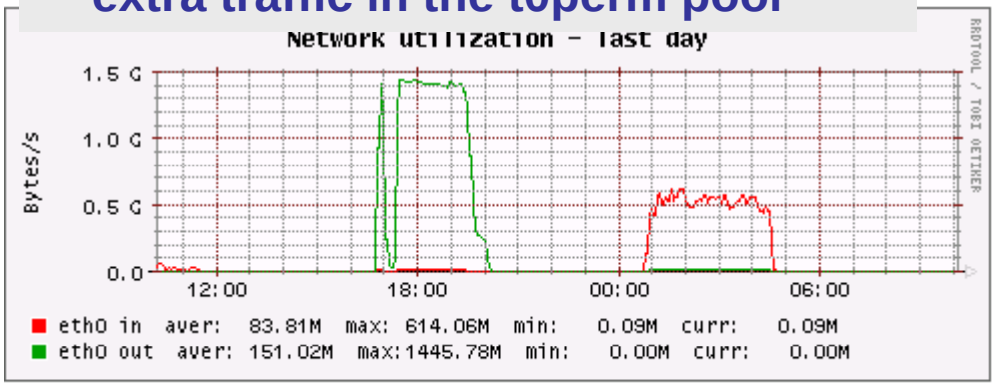

nice overlay of traffic, reaching 2.3 GB/s in and out
→no problem with the available hardware resources
== no performance bottleneck for the export

Monitoring and Visualization Tool for LCG

Data Transfer | Job Status | Service Availability

(Version: gridview-3.0.2, Installation Date: Jan 08, 2007)

<< ABOUT ABOUT >>

Current Status

(Data Transfer For 'Atlas' From All Sites To All Sites)

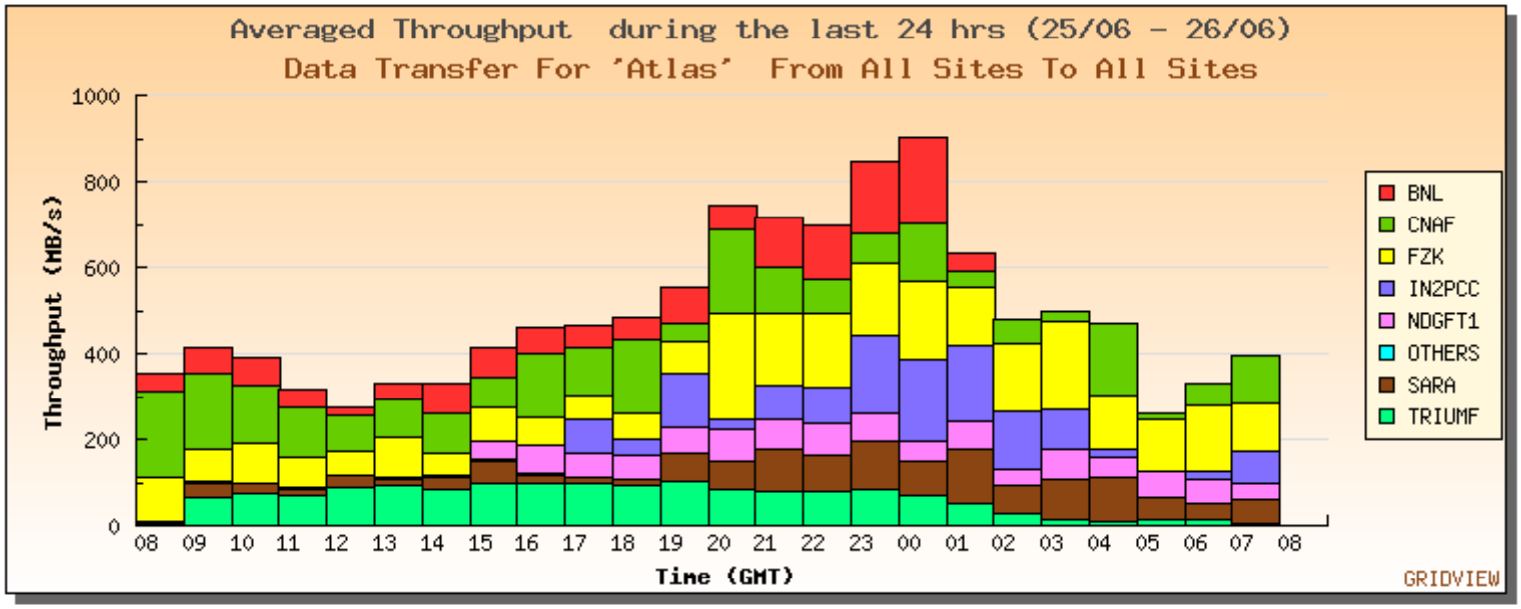

(**OTHERS:** Sites giving throughput less than 1% of max, click here for names)

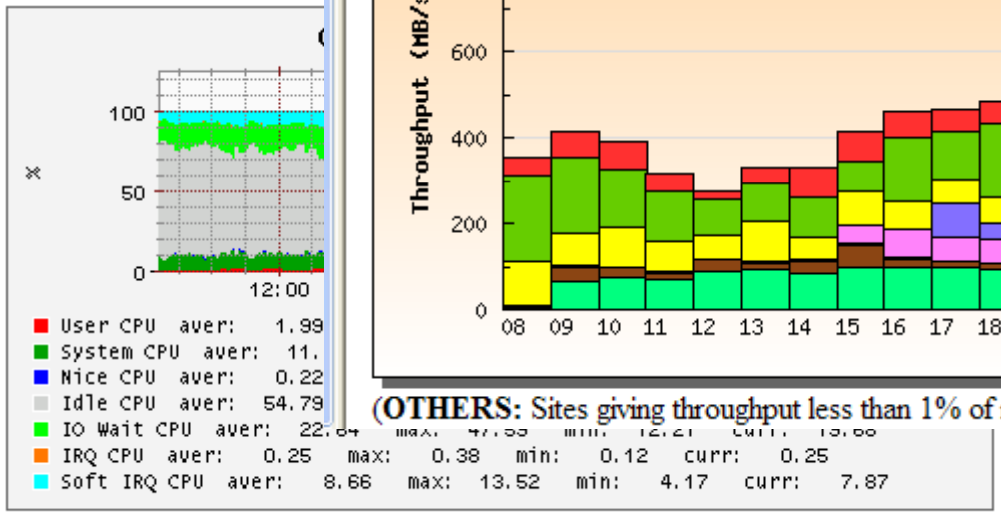

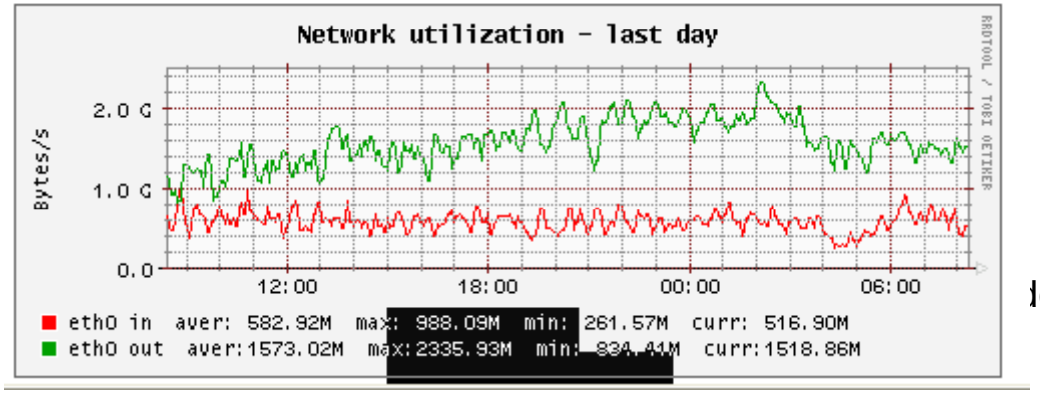

- no performance bottleneck (disk or network) or Castor scheduling problem
- up to 900 MB/s export data rates (to first order…)

Load-balancing very important

- much better than before, now externalized function, „easy“ to adjust
- needs more tuning and understanding, medium term activity
- complex area which depends on a lot of parameters
  - → number of streams
  - → number and type of disk server
  - → IO characteristics of streams
  - → TCP parameters for the export
  - →Linux IO and file system specifics
  - → ……

3d plot of the network performance for 45 disk servers over a period of 24h running the ATLAS T0 exercise

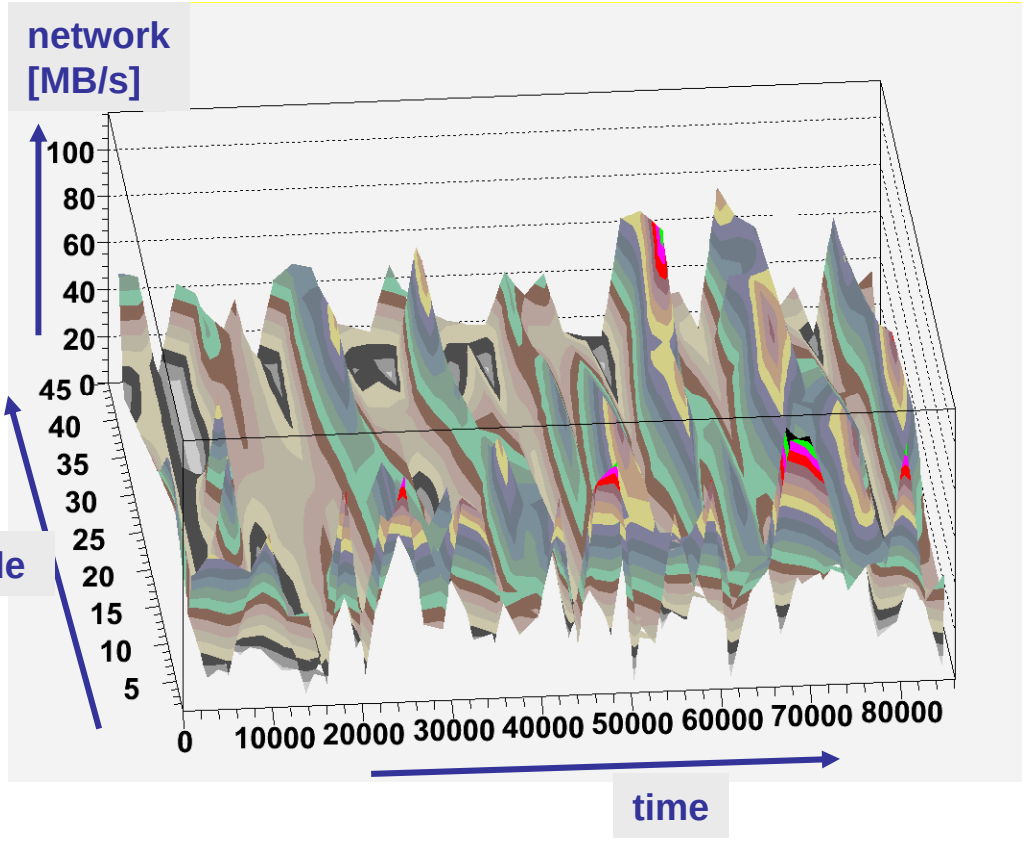

…..and also increase the monitoring effort

we have now good confidence in this release
-- used successfully by ATLAS
-- successful stress tests
-- fixes a large number of know bugs and issues

very good and very hard work from the Castor team !

during the last 3 weeks we have moved
ATLAS, ALICE, PUBLIC, LHCb and CMS
to the new Castor version.

we can now also start to concentrate again on the SRM developments and test

the focus is still on stability

Schedule I

- Restart the SRM tests and developments,
  integrate SRM test suite into the ‘standard’ Castor test suite

- The xrootd-Castor interface looks promising, passed last week
  successfully several stress tests
  discussion about deployment details will start ASAP with ALICE

- Concentrate on the Export issues in ATLAS

- More large scale stress tests to establish further the limits of the new Castor release (next ~2 month), improve the load-balancing and monitoring; improve the „chaos-containment“ procedures; some identified weakness in the tape area

- Preparation for the upcoming experiment exercises:
  - pre-CSA07 in the second part of July for CMS
  - Muon runs in August/September for ATLAS
  - Current DAQ-T0 test for LHCb
  - Integration of ‘real’ DAQ setup for ATLAS and CMS in September

Schedule II

- preparing analysis scenarios for large scale tests
- Based on the test results a general judgment of Castor has to be done, i.e. where are architectural issues which need a redesign on the 12-18 month time scale (next ~3 month)
- Arrange ASAP the deployment of the new Castor release on the external sites (RAL, CNAF, ASGC)
- New Castor version (2.1.4) planned for August ('durable' functionality, further operational improvements, etc.)

very good start but much more work to be done