

# Castor Task Force Status

The mandate of the Castor Task Force:

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

Team leader is Bernd Panzer-Steindel

The team includes : Castor operations , Castor development, members of GD and DES group plus close collaboration with the ATLAS operation team

The activity started fully in the first week of April.

- ❑ During the last 2 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
- ❑ several dedicated longer coordination and planning meetings
- ❑ high load on the team as they had in addition to support the existing 4 different Castor production releases and deal with their regular deficiencies.

started with release 2.1.3-3 → production release is today 2.1.3-14

- ❑ ~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 4 weeks**

major changes:

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

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


→ export on average 500-600MB/s with short peaks  
at 800 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



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

**3 disk pools  
(disk1tape1 and disk1tape0)**


## 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

## 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
- the upgrade procedures were tested on two different Castor instances (pre-production and ITDC data challenge)
- fixes a large number of know bugs and issues

**very good and very hard work from the Castor team !**

we have an increasing number of incidents on the current Castor production instances, which stresses the operations team considerably.

we have the upcoming variety of tests from the LHC experiments (pre-CSA07, FDR, etc.)

**→ would like to deploy ASAP the new release in production**

## Tentative schedule :

1. deploy the new Castor production release for ATLAS, PUBLIC, LHCb and CMS during the next 4 weeks at CERN;  
ALICE later, coupled to the Castor-xrootd deployment discussion
2. 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;  
understand and improve the current ATLAS export performance  
(integrate the SRM test suites)
3. 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)
4. arrange ASAP the deployment of the new Castor release on the external sites (RAL, CNAF, ASGC)