

# CASTOR and EOS at CERN

Massimo Lamanna

IT/DSS  
Filesystem and Data Operations

# Status

- CASTOR and EOS
  - Now ~18 PB usable (#replica=2)
  - Introduction of EOS side-by-side to CASTOR
  - CASTOR: Tape backed up activities (T0 function)
  - EOS: User analysis (D1)
- EOS: User analysis (D1T0)
  - Two instances in production (ATLAS and CMS) for a total of ~5 PB of usable disk
    - ALICE in the pipeline
    - Informal interest in LHCb
    - No plans for “Public”
- Streamlining of CASTOR (D0T1)
  - Dependability and manageability
    - Summer 2011: *Transfer Manager* (replaces LSF; stability/simplification)
    - Now: *Tape Gateway* (Tape performance and stability)
    - March 2012: CASTOR version 12 (2.1.12.x)

# Good examples

- CASTOR/EOS ATLAS
  - After extended joined tests: disk activities from CASTOR to EOS
  - CASTOR:
    - April 2011: ~5 PB disk space across 2 instances (ATLAS+CERNT3), 14 pools (service classes, subcluster), 400+ servers
    - January 2012: ~3 PB disk space on 1 instance with 7 pools 223 servers.
      - 2 more pools will be retired in the next couple of months.
  - Capacity moving to EOS (“user” pools)
  - Similar trend in CMS
- CASTOR LHCb
  - Restructuring of pools (service classes)
  - Neat structure built around two CASTOR pools: LHCBTAPE (D0T1) and LHCBDISK (D1)
  - Very similar to ALICE

# Protocols

- xroot as main protocol
  - Going towards deprecating rfiio (external clients and internal traffic)
  - User “local” access (lxbatch, etc...)
    - WAN use cases possible (federation). Attractive. Operational costs?
- SRM
  - CASTOR (v 2.11): Workhorse of data distribution. Dependable.
  - EOS (BESTMAN): OK but intended more as a tool for migration than as long-term solution
  - Uncertain long-term future. Replaced by FTS + gFTP? Other protocols and clients? More client responsibilities? Changing the model will change the players in the area of managed transfers as well (data placement). Support and operations?