

Tier0 Status

Tony Cass

LCG-LHCC Referees Meeting
18th November 2008

- Resource Ramp-up
- CASTOR status and performance
- Progress with new data centre project




- Resource Ramp-up
- CASTOR status and performance
- Progress with new data centre project



- CPU
 - First batch
 - Ordered out in late August
 - Delivery before November **2 of 3 batches already on site**
 - Production in December or early 2009
 - Second batch
 - Received the tender answers
 - Target FC approval in December **FC approval not required, but delivery schedule unchanged (installation depends on readiness of racks)**
 - Delivery before March 2009
 - Production in March - April 2009
- Disk
 - First batch
 - FC approval last week
 - Delivery in ~~December~~ **January**
 - Production ~~January~~ 2009 **February**
 - Second batch
 - Received the tender answers
 - Target FC approval in December
 - Delivery before March 2009
 - Production in March - April 2009
- Tape
 - Media availability not a problem but exact procurement schedule depends on progress with new repack service but **70 Sun T10KB drives ordered (1TB/cartridge) T10KA drives to be phased out as repack advances.**


Network utilization




■ eth0 in aver: 316.2M max: 933.8M min: 4.2k curr: 22.9M
■ eth0 out aver: 370.9M max: 822.3M min: 3.7k curr: 71.0M

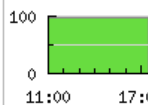
CASTOR Tape REPACK

TAPE REPACK: Completed tape repacks in last 24hrs - last month





 how is availab
 REPACK Perfor

 availability in


 11:00 17:00

Admin
admin tools

SLS by CERN IT

- Resource Ramp-up
- CASTOR status and performance
- Progress with new data centre project

- Resource Ramp-up
- CASTOR status and performance
 - Upstream services (SRM, FTS)
 - CASTOR status & plans
- Progress with new data centre project

- SRM 2.7 release is delayed
 - Originally foreseen in June but has still not yet passed testing/certification
 - Continue with 1.3
 - SLC3 – hardware running out of warranty → retire/replace

*Pre-production clusters in service for all LHC VOs
Production deployment before end-2008*


- FTS 2.1 passed certification too close to LHC startup

- Continue with 2.0 service (SLC3)
- Setting up an independent 2.1 production service (SLC4) in parallel allowing VOs to move when convenient

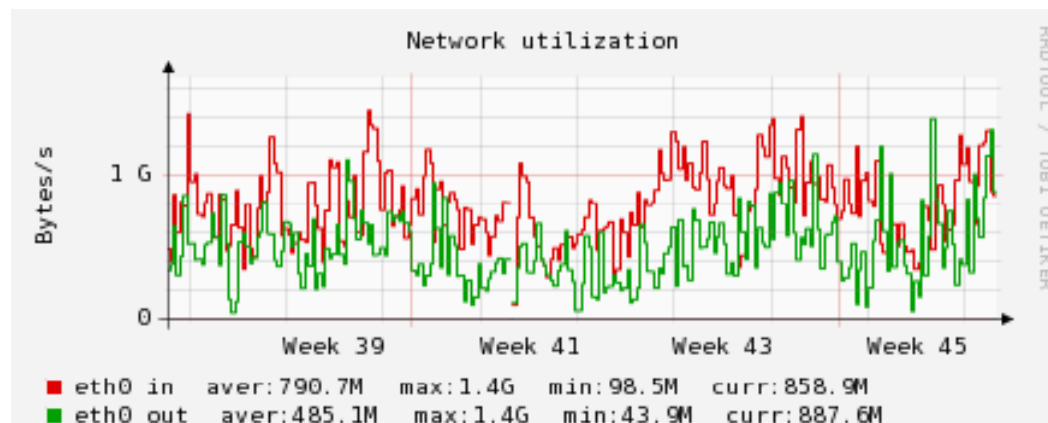
*FTS 2.1 production service available
Still being "tested" by experiments but most
production transfers already with this version*

- No accelerator... still there is data moving

- Disk Layer



- Tape Layer



- Metrics have been implemented and deployed on preproduction cluster
 - Data collected in lemon
 - RRD graphs not yet implemented
- Production deployment delayed for several reasons
 - New metrics imply several changes to exception/alarms and automated actions used in production
 - An unexpected technical dependency on the late SRM 2.7 version
 - Ongoing work to back-port the implementation

All still true 😞



DM

Security, SRM and Monitoring

- Strong authentication now implemented
 - End user access to CASTOR has been secured (nameserver, disk cache, protocols)
 - Support both PKI (**Grid certificates**) and **Kerberos**
 - Consolidation ongoing, Performance optimizations planned
 - See Frohner, Garcia, Ponce at Hepix Fall 2008
<http://indico.twgrid.org/contributionDisplay.py?contribId=133&sessionId=90&confId=471>
- Castor **Monitoring** and **SRM interface** being improved
 - See De Witt, Lo Presti, Pokorski, Rekatsinas at Hepix Fall 2008
<http://indico.twgrid.org/contributionDisplay.py?contribId=171&sessionId=90&confId=471>



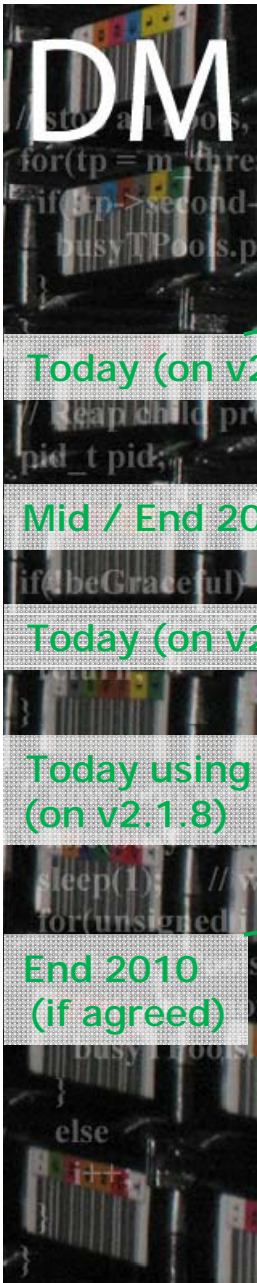
Increasing Tape Efficiency

- New tape queue management supporting recall / migration **policies**, access control lists, and user / group based **priorities**
- New tape format with no performance drop when handling **small files**
- Data **aggregation** to allow managed tape migration / recall to try to match the drive speed (increased efficiency)
- See Bessone, Cancio, Murray, Taurelli at Hepix Fall 2008
<http://indico.twgrid.org/contributionDisplay.py?contribId=181&sessionId=90&confId=471>
- See plan and timelines at
<https://twiki.cern.ch/twiki/bin/view/DataManagement/CastorTapeSubsystem>

Today (on v2.1.8)

Mid / End 2009

End 2009



Reduced access latency for analysis

- **Removed the LSF job scheduling** for reading disk files
 - Latency reduces from seconds to milliseconds (like a file system)
 - Plan to remove job scheduling also for write operation
- **Focussed to offer direct access to storage on Castor disk pools**
 - Direct access offered with the XROOT interface which is already part of the current service definition
 - Excellent integration with root
 - Mountable file system possible
 - Additional access protocols can be added (eg. NFS 4.1)
- Performances can be increased by **improving the name server**
- See A.Peters at ACAT 08

<http://indico.cern.ch/contributionDisplay.py?contribId=171&sessionId=29&confId=34666>

Today (on v2.1.8)

Mid / End 2009

Today (on v2.1.8)

Today using FUSE (on v2.1.8)

End 2010 (if agreed)



- Resource Ramp-up
- CASTOR status and performance
- Progress with new data centre project



- Reminder: the selected strategy is to do a single tender for an overall solution
- Four phase process developed:
 1. Request (many) conceptual designs
 2. Commission 3-4 companies submitting conceptual designs to develop an outline design
 3. In-house, turn a selected outline design into plans and documents enabling
 4. Single tender for overall construction.

- Deadline: 28th November
 - Contacts with all 4 companies during design phase
 - All 4 companies say deadline will be met
- Meetings to review proposed designs scheduled in week of December 8th.
- Market Survey in preparation as first stage in selection of company for detailed design & construction.
- Discussions in Oslo on 28th November to further investigate possible remote server installation in 2011 (and beyond)
 - RAL also have power available in 2011, but not as much and for a shorter period.

