FIO Fabric Infrastructure and Operations

CERN IT Department

CASTOR Status

March 19th 2007

CASTOR dev+ops teams

Presented by Germán Cancio

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Outline

- Current status/issues
  - ATLAS T0 challenge
- Current work items
  - SRM-2.2 – *see previous presentation*
  - LSF Plugin
  - New DB hardware
  - xrootd interface
- Medium term activities
  - Tape repack service
  - Common CASTOR/DPM rfio library
- Future development areas

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Current status/issues (1)

- The current production release (2.1.1-4) was deployed in Nov'06.

- Bugs and issues:
  - LSF plugin fails unrecoverable way upon high load. Workaround: automatic LSF master restart (up to 5 mins of non-service)
  - DLF (logging) client socket leak →stuck servers
  - Bug in tape error handling module. May cause insufficient retries upon failing recall/migration
  - Known performance/interference bottleneck in LSF scheduler plugin

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Current status: ATLAS T0 challenge

3-4/3: Problem with DLF DB → socket leak in two central servers

5/3:DLF problem fixed. New problem: corrupted block in stager db.
- DB shutdown and recovered
- De-fragmentation of largest table

6/3: Problem with oracle execution plan causing high row lock contention

7-11/3: DB OK but clients see high error rate due to unrecoverable errors in LSF plugin.
- Deployed a new binary without DLF

12/3: 5 more diskservers added to 'default' to better cope with non-T0 activities

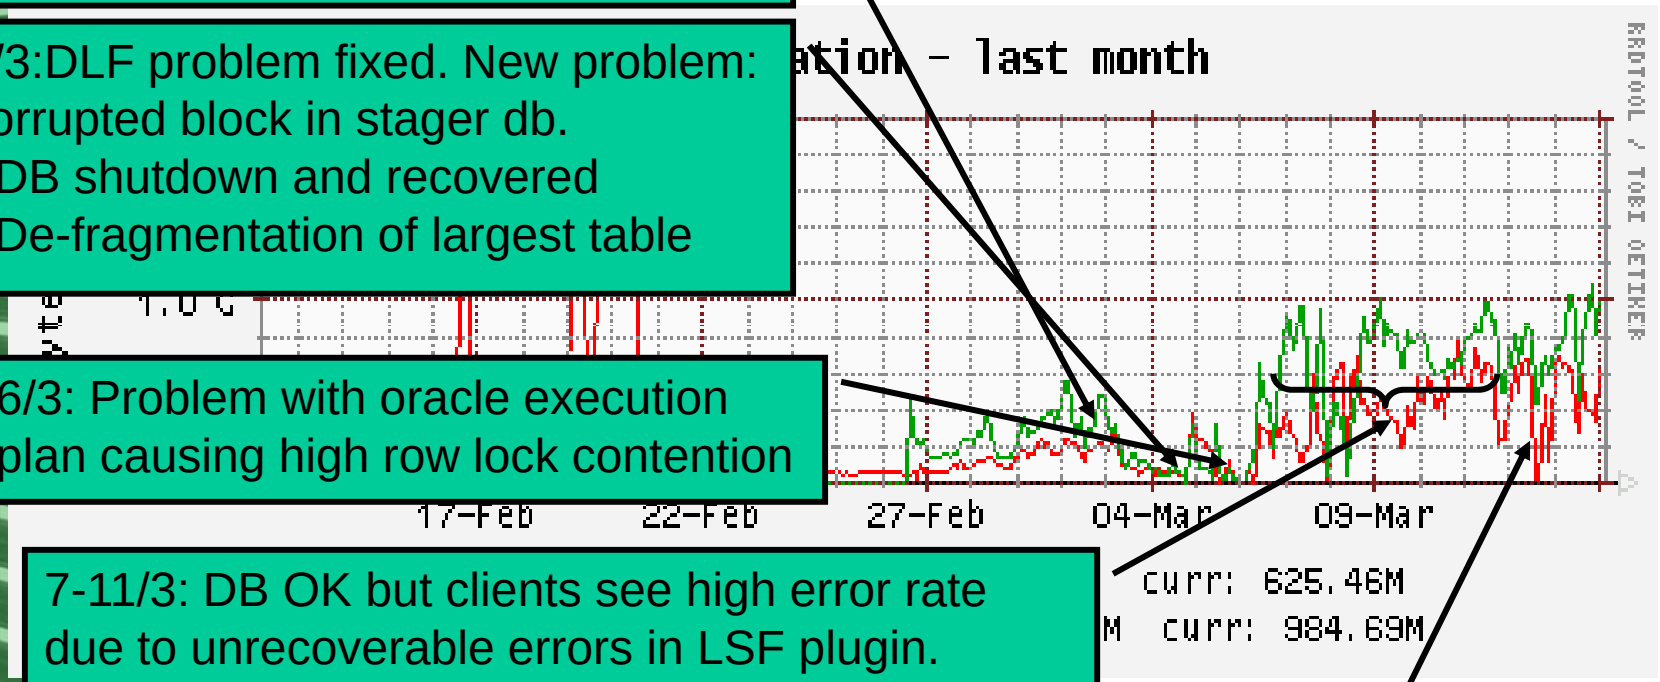

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Current status/issues (2)

- Known LSF limitations (details next slide) in current release prevent the stager from supporting loads ~>5-10 requests/s
  - This limitation also affects migration of non-LHC CASTOR1 instances (e.g. NA48, COMPASS) to CASTOR2
- Hardware/Oracle problems
  - Broken mirror errors followed by performance degradations during mirror rebuild
  - Block corruptions requiring DB recovery
- These problems are being addressed with top priority, though the timescales for deploying new releases did not match the dates for the ATLAS T0 challenge
  - A newer release with improved Oracle auto-reconnection to be installed ASAP
- Other recent developments which are approaching production deployment are
  - More efficient handling of some requests used by SRM
    - prepareToGet and putDone requests are no longer scheduled in LSF. Represents ~25% of the ATLAS request load
  - Support for the xrootd protocol in CASTOR2, which may help to shield chaotic analysis activity from the inherent latencies in CASTOR2 request handling

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

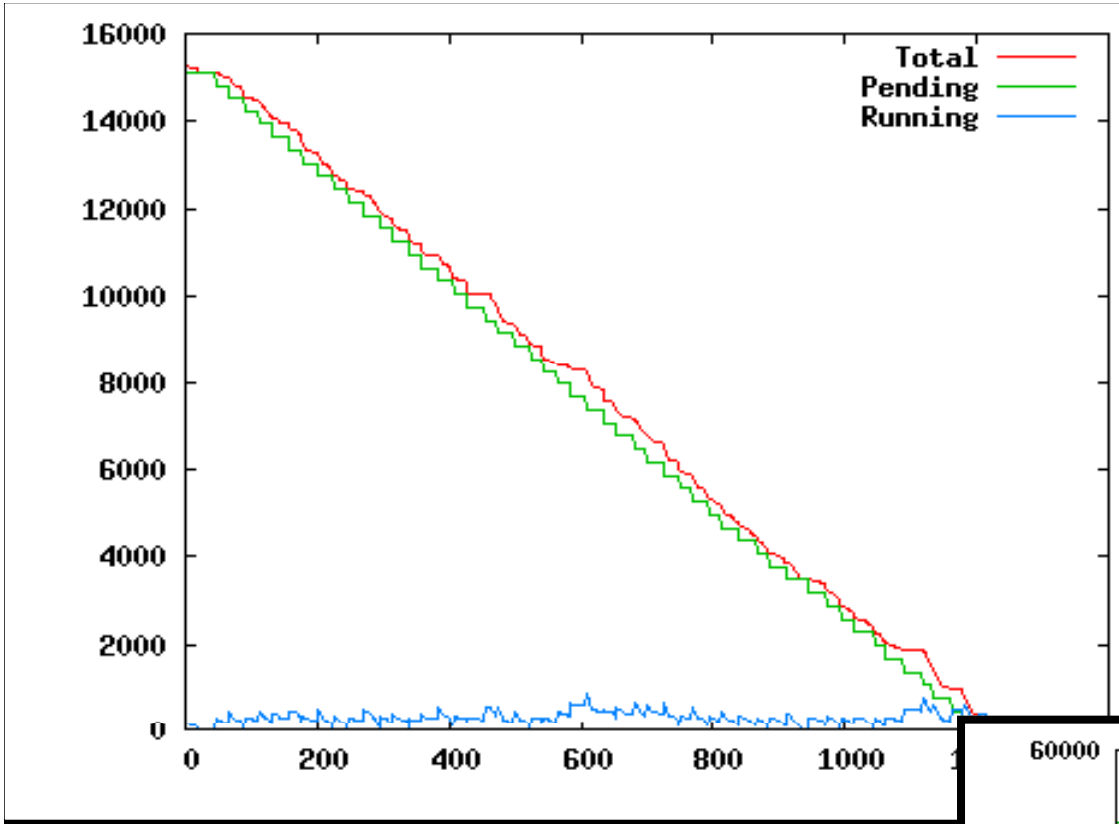

CERN IT Department

d identified as priority area during

OR specific LSF plug-in leads to execution

scheduling:

uests to LSF to schedule to minimize
dule 1,001st request according to

d when

0 slots

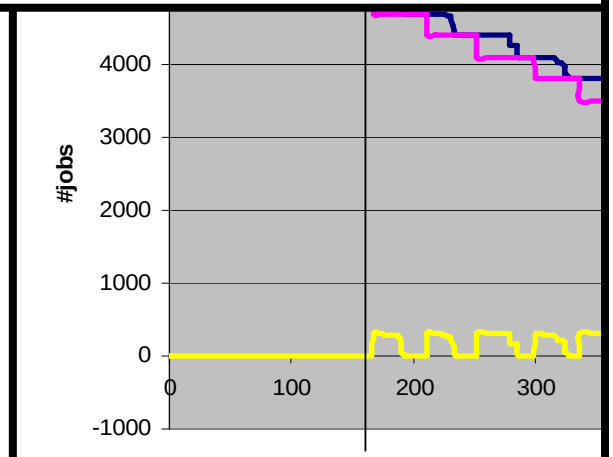

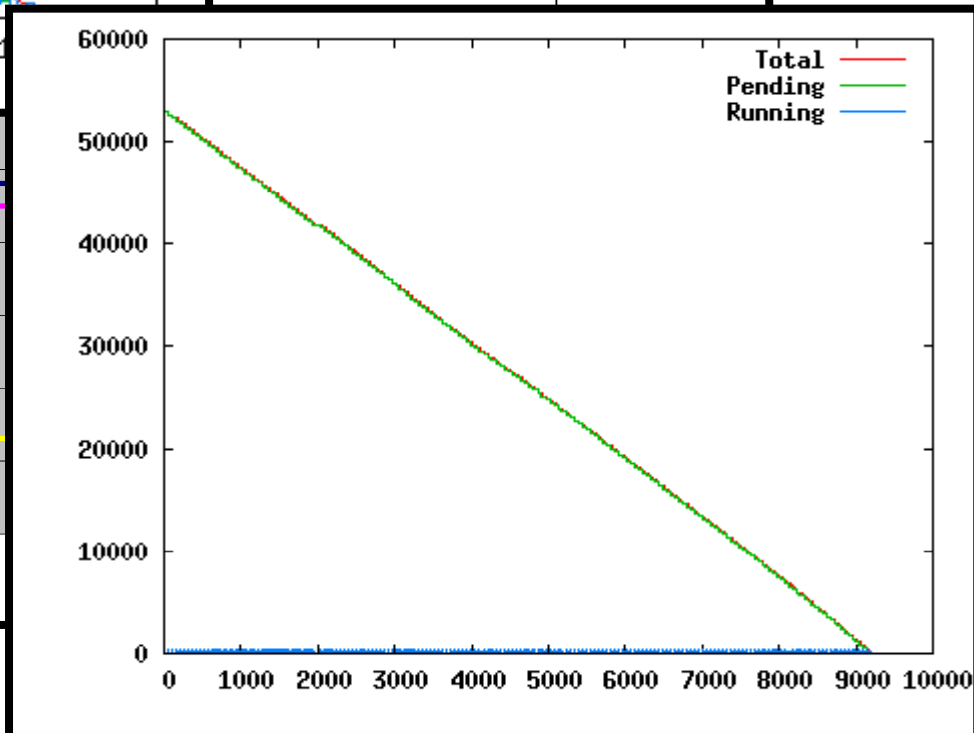

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

New DB hardware

CERN IT Department

- New database hardware:
  - RAC clusters of oracle certified hardware has been purchased for hosting the name server and stager databases
    - CASTOR name server move 2nd of April
    - Definite plans for moving the stager databases will be negotiated with the experiments (~1/2 day downtime)
  - The SRM v2.2 database will also be moved before production deployment

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Xrootd interface

- Core development done Q3 2006, testing Q1 2007 by SLAC (A. Hanushevsky) with support from CERN
- Current status
  - Dev setup already open to ALICE test users
  - CASTOR and XROOT RPMs are ready and will be made available this week on the pre-production setup.
  - ALICE instance will be upgraded as soon as verified by ALICE.

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Tape Repack

- A new tape repack component is required for CASTOR2 for several reasons
  - in the short term, to efficiently migrate 5PB of data to new tape media
  - in the long term as we expect tape repacking to become an essentially continuous background activity
  - old Repack was CASTOR-1 specific
- Current status
  - Running new version on pre-production instance for 2 months, intense debugging
  - Testing long, concurrent repacks
  - Pure Repack functionality now complete, but tests are highlighting need for stability and reliability improvements in the underlying core CASTOR2 system.
  - While working on Repack2, migration to LHC robotics from 9940 being performed using CASTOR1- O(500) days / 20 drives to complete (14K tapes left, out of 21K)

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Common CASTOR/DPM RFIO

CERN IT Department

- CASTOR and DPM have two independent set of RFIO commands and libraries, but with same naming
  - Has represented a source of link time and run time problems for experiments
- A common framework architecture allowing MSS-specific plug-ins has been designed and implemented
  - Plug-ins for CASTOR2 and DPM developed
  - common CVS repository contains framework and clients (e.g. rfcp,…)
- The new framework has been extensively tested on the CASTOR dev setups
- Coordination with DPM developers for defining the production deployment schedule is ongoing

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN

Future Development Areas (1)

CERN IT Department

- We expect development work on critical areas (especially the LSF plugin and SRM v2.2) to be complete by mid-year.
- Development plan for second half of the year is still to be agreed, but two issues are likely to be addressed with priority:

1) Server Release support: Simultaneous support requested by external institutes for two stable major server releases
- Comparable to supporting SLC3 and SLC4:
  - Critical bugfixes backported to more mature release
  - New functionality will go only to newer release
- Release lifetime defined by available support resources – aiming for O(6) months
- Client releases: backwards compatibility defined by user needs
  - Mechanism for ensuring backwards compatibility in place

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Future Development Areas (2)

2) Strong Authentication:

- Current uid/gid-based authentication is not secure
- Strong authentication libraries for CASTOR were developed in 2005 and integrated in DPM but not yet in CASTOR2
- Developments ongoing, to be completed in Q3 2007
- Rollout / Transition phase needs to be defined, in particular during lifetime of Castor1 (common name server)
- Improved ACL support for VOMS roles and groups will be built on top

CERN - IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it