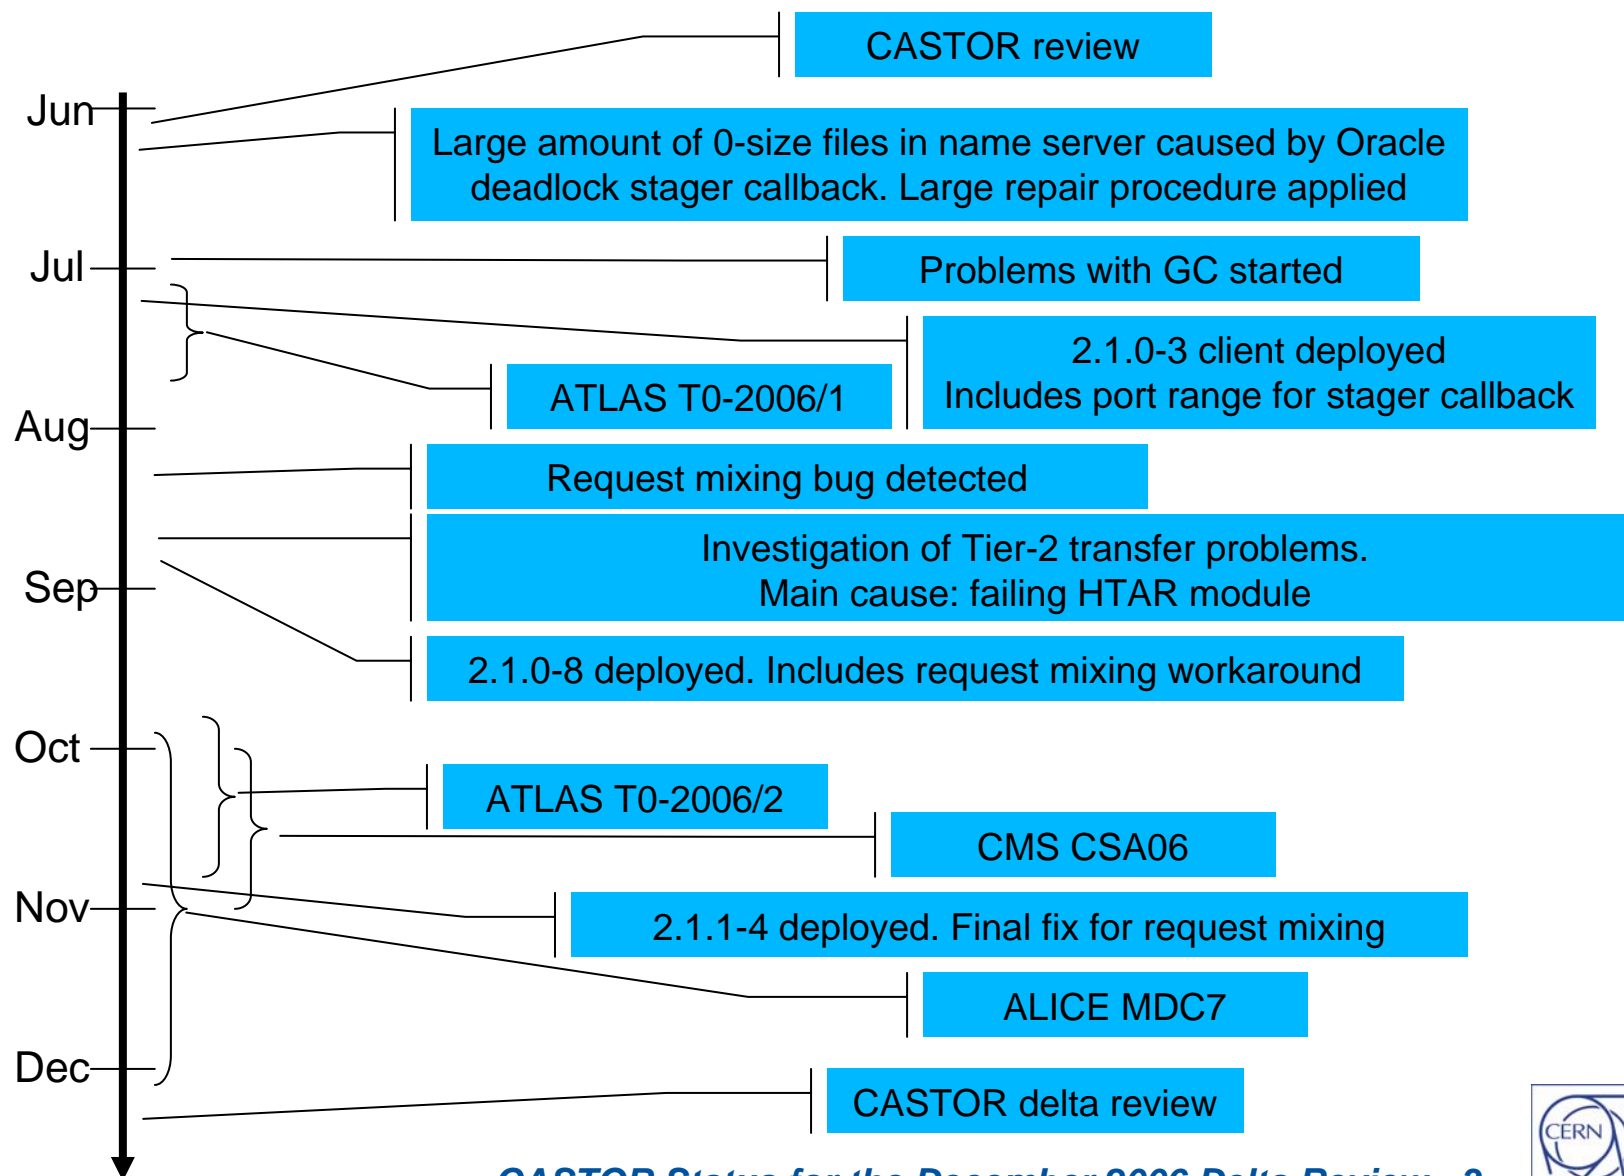


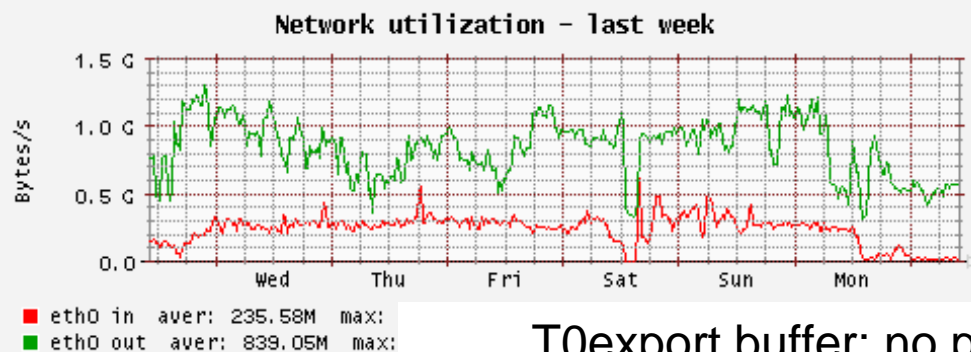
CASTOR Status

for the
Delta Review, December 6th 2006

Tony Cass


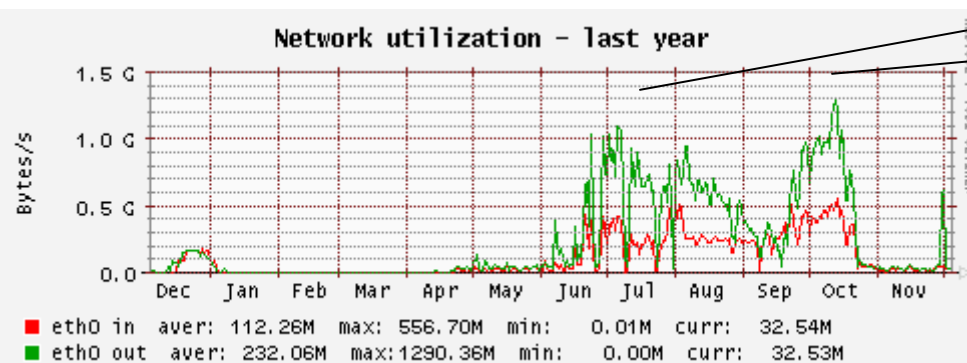
- CASTOR was late and uptake slow
 - So in red in IT DH presentation, January 2006
- Regular appearance of new problems in 1H06
 - Fixed rapidly, but sapped confidence
- Review: Hard area; but OK for the long haul
 - Repeat: for the long haul. Invest effort and don't expect miracles.
- Since then
 - Some new problems, but
 - these generally appeared for new use cases or as side effect of other changes (c.f. request mixing).
 - Also major successes
 - At CERN and elsewhere.





CMS CSA06

T0export buffer: no problem...

ALICE MDC7 target
reachedATLAS T0-2006
Phase 1
Phase 2

- Stager & client: Port range and request ID for callback
 - Fixes request mixing problem,
 - although still subtle concurrency problems at bind/listen level.
- Improved DLF (with reassigned effort)
 - Improved server side scalability plus asynchronous API to eliminate client overhead
 - Log archiving with Oracle tools
 - Extended GUI
- Improved garbage collection
- System interface improvements
 - XFS preallocation
 - O_Direct support in RFIO
- LSF plugin/monitoring rewrite
 - Major development for performance; see talk later
 - Ongoing, but release to production in near future.
- Automatic database reconnections
 - Improves robustness wrt network glitches.
- New Repack component
 - Essential for efficient migration from 9940 to T10K/3592.

- IBM 700GB tape
 - Essential: media costs dominate @ CERN
- Improved compression reporting
 - For long term efficiency generally, but also to profit from IBM non-volatile cache support for small files.
- SLC4 & 64 bit
- Software Build, Test Suite & Release Improvements
 - See talks later

- RFIO
 - Merged with DPM code base, with common CVS
 - Extracted from the CASTOR source code
 - In collaboration with DPM
- GridFTP v2
 - Integrated both as internal and external protocols
 - Already tested with SRM v2.2 on test setup
 - Problems with distribution of gridFTP (no RPMs)
- XRoot
 - Integrated as internal protocol
 - Tests still not over, especially for put case
 - Developments done by xroot team
 - Many thanks to Andy Hanushevsky

☺ Fixed

- Many database deadlocks
- Bad statuses & accumulations in the database
- Stager permission problems
- Stager_qry problems

☹ Left

- One database deadlock
- Problems with readonly file creation
- Problems with putDone on multiple files
- Bind/listen not thread safe

● Unfound

- We don't [have time to] talk about unknown unknowns!

- Installations at CNAF and RAL now ~OK.
- Discussions on External support
 - Clarified that CERN is 3rd level support, not 2nd level operational support for sites.
 - Need for experienced staffing emphasised
 - Would also be needed for other MSS products
 - Even if we admit CASTOR robustness is not at the level of commercial products
 - CNAF effort is ramping up significantly.
- “Disk 1” support
 - Also an issue for CERN; but if we can use CASTOR for this then other sites should be able to as well
 - But be realistic: 9x% disk occupancy rates are not going to be possible.
- IHEP using Oracle XE & Maui

- Team ~stable and January recruits now well integrated.
 - Embed in FIO/FD has allowed some staff flexibility (e.g. for DLF) and more likely in 2007
 - Although this is a two edged sword...
- Planning externalised (FIO/FD SL)
 - But GANTT planning is still optimistic
 - More issues than we'd like, even if we take this rule into account.
 - Unrealistic expectations of team's ability to deliver everything requested when requested
 - c.f. AuthN & AuthZ. High priority for 2007.
- Team cannot deliver Windows & Mac clients
 - We don't have the skills.
 - But those that need these clients do, and have delivered working code in the past.
 - We would *really* like them to work with us now to deliver integrated code, not just code that works one off.

