



Dissemination and User Feedback

Castor deployment team



Dissemination



- ❖ **Conference contributions** (*selected from Castor website*)
 - **CASTOR status and overview**
HEPIX, Apr 2006
 - **CASTOR: Operational issues and new Developments**
Operation of the CERN Managed Storage environment
CHEP, Sept 2004
 - **CASTOR/New stager presentation,**
Storage Resource Sharing with CASTOR
NASA/IEEE Mass Storage Conference, Apr 2004
- ❖ **Regular reporting**
 - PEB, GDB, LHCC, IT department
- ❖ **Training**
 - **CASTOR2 operation training** *Dec 2005*
 - **External operations workshops** *Feb 04, Nov 04, Jun 05, Jan 06*
- ❖ **Documentation, user guides, admin guides, man pages....**
 - lots of it
 - various web sites, reflecting organizational updates
 - in urgent need of restructuring!



Migrate LHC expts to Castor-2



- ❖ Main activity since late 2005
- ❖ Questionnaire, discussions to configure according to needs of experiments
- ❖ Identify missing functionality, and provide it 😊
 - enhanced `stager_qry` command
- ❖ Agree steps for migration:
 - initial configuration (and sizes!) of diskpools
 - staged migrations: production, data export, user analysis
 - stop usage of `castorgridsc.cern.ch` (SC3!)
- ❖ Migrate experiments one-by-one...
 - but as quickly as possible



General remarks



- ❖ *From our perspective* 😊
- ❖ More similarities than differences between experimental needs (fortunately)
 - Castor accessed through experiment frameworks
 - `stager_qry` enhancements were required by all
- ❖ Discussions and presentations helped a lot to avoid confusion, and prevent (many, but not all...) problems
- ❖ Migration provided good opportunity to clean up
 - disk caches flushed!
 - “temporary” SRM mappings...
- ❖ Confusion about required and provided connectivity resulted in many tickets 😞



General remarks (2)



- ❖ Challenging to communicate effectively with Offline coordinators, production managers, end users, T2 managers
- ❖ Support line Castor2.Support@cern.ch well known by now
 - advertised to all experiments and their users, also part of Grid Support lines
 - it is being restructured on our side:
 - 1st level: FIO Service Manager on Duty (*rota*)
 - 2nd level: Castor Service Manager
 - 3rd level: Developers



Alice



- ❖ discussions with Latchezar Betev
- ❖ well understood usage of Castor ☺
- ❖ first to go, presented on February 2
- ❖ additional complication: xrootd
 - was running on (public!) `castorgrid` nodes
 - is now running on dedicated Alice machines, with special configurations...
 - connectivity required for transports to non-Tier1's
- ❖ Next: Alice data challenge in July



LHCb



- ❖ discussed with Philippe Charpentier, Joel Closier
- ❖ presented on March 8
- ❖ require rootd on diskservers
 - found several problems, fixed by Root and Castor teams
- ❖ require 'durable' SRM endpoint
 - Castor release 2.1.0, to be deployed soon
- ❖ user migration still ongoing
- ❖ only experiment to ask for Windows client (*for online needs*)



CMS



- ❖ discussions with Nick Sinanis
- ❖ presented on March 9
- ❖ 2 independent Castor-1 stagers
 - stagecmsprod
 - stagecms
- ❖ quite some “SC3 activity” until March/April
- ❖ migration completed beginning of May
- ❖ problems:
 - connectivity required for PHeDex transports
 - load on diskservers caused by normal(!) user activity



Atlas



- ❖ discussions with Gilbert Poulard, Luc Goossens
- ❖ presented on March 27
- ❖ most affected by Castor-1 limitations
 - very frequent crashes
- ❖ require rootd on diskservers (*like LHCb*)
- ❖ require 'durable' SRM endpoint (*like LHCb*)
- ❖ migration completed end of April
- ❖ problems
 - failing hardware for stager database
 - after that: stager database problems...
 - load on diskservers caused by normal(!) user activity
- ❖ we are now preparing the TDAQ data challenge



Service Challenges



- ❖ very demanding (and useful!) customer!
- ❖ simple use of Castor...
 - single diskpool, with disk resident files, read-only, large files
- ❖ ... but large setups...
 - SC3
 - service phase: many imbalanced (old) servers, complicated network topology/castor setup, handcrafted configurations, etc
 - rerun: 16 identical servers on single IP service
 - SC4
 - throughput phase: >50 servers on 4 IP services
 - service phase: ~40 servers
- ❖ ... with challenging requirements
 - SRM endpoints *to be registered in GOCDB and site BDII...*
 - SLC4
 - gridftp performance monitoring using RGMA
 - management of host certificates *still lots of handwork*
 - frequent support requests 24 x 7 ☺



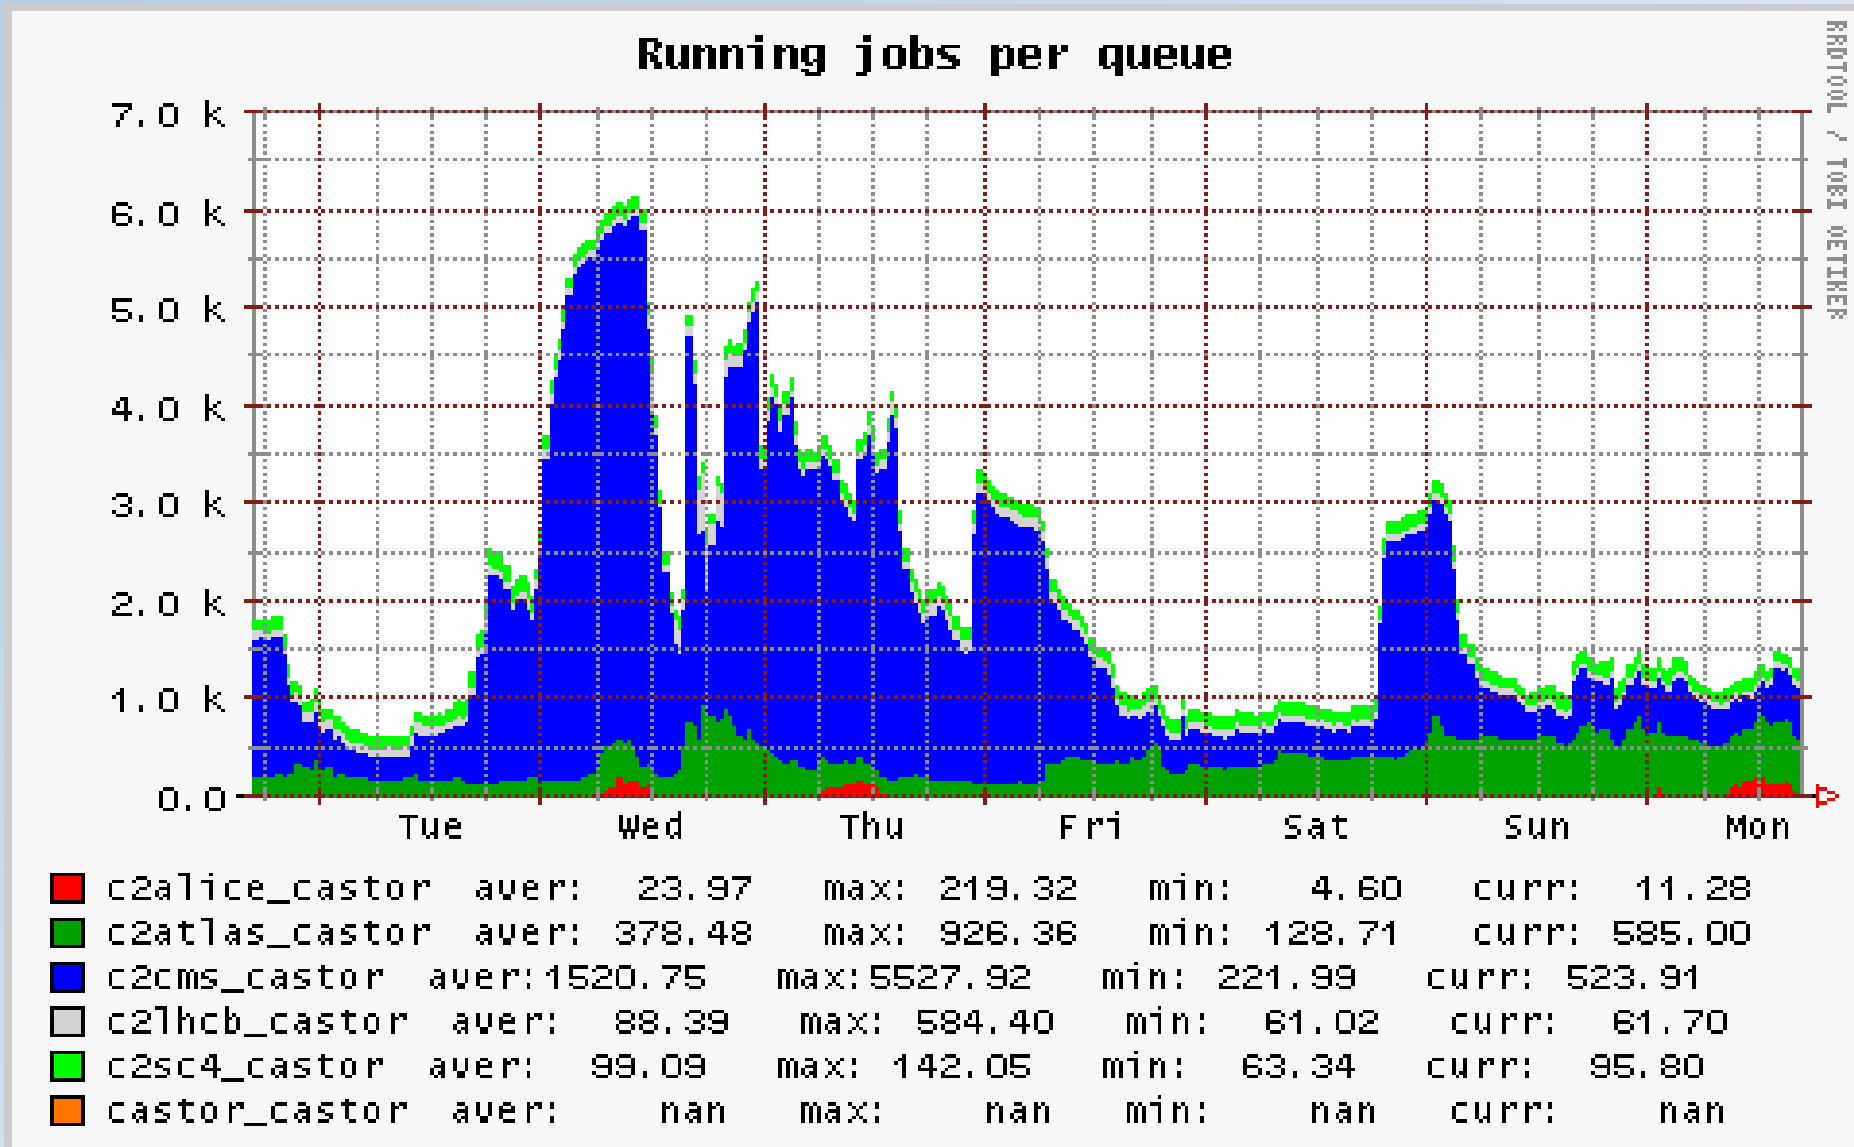
ITDC



- ❖ IT Data Challenges – *Bernd Panzer*
- ❖ Simulate (and verify) Tier-0 activities
 - many concurrent read + write streams
- ❖ Pushes Castor to its limits (*and beyond*)
 - scheduling and migration policies
 - new tape infrastructure
 - test deployment of new software releases
- ❖ Large setup, with many servers
- ❖ Lots of feedback...



Castor-2 usage: view from LSF





Conclusion



- ❖ Migration of LHC experiments almost completed...
 - thanks to the experiments for fruitful discussions
 - most problems found along the way are fixed
 - ... ***now their setups have to evolve!***
- ❖ Service Challenges and ITDC very useful
 - requirements, bug reports, field testing, etc.
- ❖ Support structure improvements
 - documentation needs overhaul
 - support flows being rationalized