



Castor-dev planning and resources for 2007

Castor Development Team

German Cancio, Giuseppe Lo Presti, Sebastien Ponce
CERN / IT

Castor Delta Review –December 2006



Outline



- ❖ Evolution of development resources since 6-2006
- ❖ Update on development plans for (Q1/2) 2007



Evolution of Resources since 06/06



Developer	From	To	%	Main development responsibilities
Olof Barring (FIO/FS)	1999	now	5	Castor-1, tape area maintenance, bughunting
Hugo Cacote (FIO/TSI)	11/05	now	60	Tape driver software (stepping down)
Felix Ehm (tech student)	11/05	12/06	100	Repack
Rosa Garcia Rioja	03/06	now	100	GridFTP/xrootd, 64-bit port, web pages, strong authentication
Giuseppe Lo Presti	03/05	01/08	100	SRM-2, stager + DB code, ETICS integration
Sebastian Lopienski	11/05	06/06	50	Strong authentication
Sebastien Ponce	09/03	now	100	Project leader+main architect, stager, scheduler
Miroslav Siket	01/07	now	30	Tape system optimisations (file recalls)
Giulia Taurelli	01/06	now	100	Test suite, RFIO, stager
Dennis Waldron	11/05	now	50	DLF
Arne Wiebalck (FIO/LA)	11/06	now	50	Tape driver software (taking over from Hugo)
Victor Kotlyar, Vitali Motyakov (Russia)	07/03	now	0.5 FTE	DLF, GridFTP-2
Shaun de Witt (RAL)	04/05	now		

Few changes since Q2 2006 (compared to Q4 05 -> Q2 06):

- Team members new in Q1/Q2 06 now trained and up-to-speed
- S. Lopienski reassigned to other tasks outside CASTOR
- H. Cacote being replaced by A. Wiebalck (in FIO/TSI)
- M. Siket (30%) from Q1 2007 on
- Technical Student to replace Felix in Q1/2 2007 (TBC)
- *Giuseppe's contract running out beginning of 2008*

- ☐ "senior" team member ☐ former team member
- ☐ "junior" team member
(≤1 year) ☐ ext. collaborator



Planning document update



- ❖ Latest version of planning document updated beginning of December
 - Now 33 pages, 66 tasks defined
- ❖ Takes into account collected tracking information (June-November)
 - Added 'ongoing' tasks such as 3rd-level support, releases, management
 - More realistic %FTE allocation for development tasks
 - Leading to postponed tasks...
 - Most tasks have now allocated manpower
 - Avoiding concentrations on Giuseppe/Sebastien
 - Spent manpower now accounted separately from progress/duration
- ❖ <http://cern.ch/castor/docs/planning>

Castor Planning	
Castor development team	
December 5, 2006	
Contents	
TG1 Known bugs	4
T1.1 Memory leak	4
T1.2 Bad states in the DB	4
T1.3 Memory leak follow-up	5
T1.4 Databases Cleanup	5
TG2 Testing	6
T2.1 Test recovery after DLF restart	6
T2.2 Test DB hardware configurations	7
T2.3 New DLF testing and deployment	7
TG3 Stager	8
T3.1 Signal handling on the client part	8
T3.2 Improvements on the stager_query	8
T3.3 Automatic cleanup at restart	9
T3.4 SvcClass specific selection policy for filesystems	9
T3.5 Recall policies	10
T3.6 Filesystem GC trigger optimization	10
T3.7 Handle recovery after DB restart	10
T3.8 Accounting	11
TG4 Scheduler	12
T4.1 LSF limitations	12
T4.2 LSF improvements	12
T4.3 Avoid scheduling prepareTo requests	13



R16 : Credible WBS (VI)



❖ Pending/ongoing:

- Define more formal (but lightweight..) mechanisms for reviewing development WBS and progress with Castor Ops team
- Identify a light-weight effort tracking tool (for both Castor-dev and ops teams) which can be used at a (bi-)weekly basis
- Include external collaborators in tracking (outside the section)
- Move to more deliverable-based rather than activity-based planning e.g. GDPM


❖ Concerns:

- High level of interruptions to Castor-development team
 - In particular for 'senior' team members
 - Significantly improved with knowledge build-up on Castor-Ops side, but still too high (e.g. Tier-1 sites)
 - A weekly support rota for developers has been put in place
- Many external dependencies make planning ahead difficult
- Management overhead not always appreciated
 - Keep it to the required minimum
 - Resources are limited and not easy to reallocate to this task
 - Not all developers see clear benefits

yesterday's presentation



Planning support tool





Planning support tool



- ❖ Still using GanttProject for helping in the definition of tasks, and their visual representation
 - chart information automatically generated from WBS document
- ❖ Q4 2006 -> Q4 2007 currently looks as follows:
- ❖ Limitations:
 - Unflexible allocation of resources to tasks
 - not possible to vary participation of a developer to a given task for a given time interval
 - Leading to occasional over-allocations
 - Time buffers look like gaps in planning
 - Task, not milestone oriented
 - Much detail required on allocations and availability (e.g. holidays, weekends, etc) which is too heavy
 - ganttproject sometimes crashing
 - WBS-to-ganttproject conversion script needs improvements..



Update on new+open tasks (1)



- TG1, TG2 (known bugs and testing)

Task	Start;duration progress	Who (%)
T1.3 Memory leak follow-up: verify if the suspected Oracle memleak is really gone so that stager cron restart can be removed <i>new – followup from T1.1</i>	15/12/06; 1d	Sebastien(25%)
T1.4 Databases Cleanup provide a PL/SQL script for cleaning up Castor DB's having old leftover stale requests <i>new – followup from T1.2</i>	1/12/06; 2d	Sebastien(10%)
T2.2 Test DB hardware configurations for stager (including RAC and DataGuard) <i>no progress</i>	Q1 2007; 5d	?

□ medium, □ high, □ low importance



Update on new+open tasks (2)



- TG3 (Stager)

Task	Start;duration progress	Who (%)
T3.1 Signal handling on client part (gracefully handle signals like SIGINT) <i>rescheduled</i>	15/4/07; 5d	Giulia (50%)
T3.2 stager_query improvements: reduce load on stager by minimising query overhead (e.g. due to regexps) <i>partly postponed</i>	1/4/07; 3d 50%	Sebastien(25%)
T3.3 Automatic cleanup at restart: remove stale DB entries after crashes of components <i>rescheduled – awaiting higher stability</i>	1/6/07; 10d	Giulia (50%)
T3.4 SvcClass specific selection policy: allow different filesystem selection policies per service classes and disk pools <i>rescheduled</i>	1/5/07; 5d	Dennis (45%)
T3.5 Recall policies: allow for a smarter scheduling and grouping of recall operations <i>rescheduled</i>	8/1/07; 5d	Dennis (40%)
T3.8 Accounting: Modify the RH to protect the stager from (DoS) request floods <i>rescheduled after CMS workaround</i>	1/9/07; 10d	Rosa (50%)

□ medium, □ high, □ low importance



Update on new+open tasks (3)



- TG4 (Scheduler)

Task	Start;duration progress	Who (%)
T4.1 LSF limitations: study and improve limitations in LSF plugin for higher job/s throughput rates <i>work in progress</i>	15/8/06; 20d 50%	Sebastien(25%)
T4.2 LSF improvements: improve error reporting, configurable # of retries for scheduling a job, study using several LSF queues with different limits <i>new</i>	8/1/07; 8d	Sebastien(25%)
T4.3 Avoid scheduling prepareTo requests: bypass LSF when launching prepareTo requests (no scheduling decision needed) <i>new</i>	1/12/06; 5d 40%	Felix (80%)

□ medium, □ high, □ low importance



Update on new+open tasks (4)



- TG5, TG6 (Protocols and Security)

Task	Start;duration progress	Who (%)
T5.1 RFIO common project: provide CASTOR plugin for new modular RFIO developed by IT/GD; adapt CASTOR code base <i>work in progress</i>	1/8/06; 45d 75%	Giulia(40%), Giuseppe(20%)
T5.2 GridFTP: add support for GridFTP2 as internal+external CASTOR-2 protocol <i>work in progress</i>	1/9/06; 40d 90%	Victor(90%),Rosa (60%)
T5.3 XROOTD: assist xrootd developers in integrating/testing support for XROOTD as internal CASTOR-2 protocol <i>work in progress</i>	1/4/06; 20d 90%	Rosa(20%)
T6.1 Strong authentication: Finish the implementation of strong authentication in CASTOR-2; define a rollout plan <i>rescheduled</i>	8/1/07; 40d 5%	Rosa(30%)
T6.2 VOMS support: Add VOMS roles/groups for authorization purposes. <i>new</i>	1/3/07; 25d	Rosa(15%)

□ medium, □ high, □ low importance



Update on new+open tasks (5)



- TG7, TG8 (Components to improve/change/rewrite)

Task	Start;duration progress	Who (%)
T7.2 DLF archiving GUI: enhance DLF GUI for managing DLF archives <i>postponed</i>	1/10/07; 10d	Dennis(20%)
T7.3 SRM-2 updates, deployment: modify existing SRM-2 implementation as agreed by WLCG for post-SC4 <i>work in progress</i>	1/5/06; 40d 80%	Shaun(30%), Giuseppe(4%)
T7.4 New VDQM: finish development, test and deploy new VDQM <i>postponed</i>	8/1/07; 30d 10%	Giulia(50%)
T8.1 CLIPS to Perl transition: replace complex CLIPS-based expert system by a simpler perl script <i>postponed</i>	1/3/07; 2d	Sebastien(20%)
T8.2 Job Dispatching: Redesign scheduler interface part of rmmaster	1/2/07; 20d	Dennis(40%)
T8.3 CUPV: redesign old/unmaintained user privilege validation module, reduce complexity and simplify authorization setup, consider using VOMS	1/4/07; 2m	? (Technical student)

□ medium, □ high, □ low importance



Update on new+open tasks (6)



- TG9, TG10 (Tools, platforms)

Task	Start;duration progress	Who (%)
T9.1 Repack for CASTOR-2: rewrite tape repacking tool using new CASTOR-2 API's, address enhancement requests by castor-ops team <i>work in progress</i>	1/2/06; 136d 95%	Felix(85%)
T9.2 Makefiles and packaging: Provide a clean, modular, well-organised build/packaging schema <i>work in progress</i>	15/7/06; 50d 15%	Giuseppe(30%)
T9.3 Automation of releases and tests: automatic regular builds and execution of test suites using ETICS framework <i>postponed</i>	15/4/06; 20d	Giuseppe(20%)
T9.4 Recover tool for Repack2: Develop an easy-to-use tool for re-activating a repacked tape (in case repack went wrong) <i>new</i>	22/11/06; 8d 50%	Giulia(40%)
T10.1 Port of servers to 64bits: port CASTOR-2 server parts to 64bit (stager, scheduler, SRM, etc) <i>work in progress</i>	1/9/06; 40d 65%	Rosa(20%)

□ medium, □ high, □ low importance



Update on new+open tasks (7)



- TG11,TG12 (Documentation, Tape optimizations)

Task	Start;duration progress	Who (%)
T11.3 Guides: Write a developer guide	1/10/07; 15d	?
T11.4 Articles: publish articles on CASTOR-2 (architecture, deployment, HA, ..) <i>postponed</i>	8/1/07; 10d	Giuseppe(25%)
T12.1 Optimization of file recalls: investigate how to optimize tape recalls (e.g. ordering of file recalls within a tape, use/tune recall policies, etc) <i>new</i>	8/1/07; 20d	Miroslav(30%)
T12.2 Support for small files: investigate ways of grouping small files for optimized tape usage <i>new</i>	1/5/07; 30d	Miroslav(30%)

medium, high, low importance



Update on new+open tasks (8)



- Ongoing activities (all TG)

Task	Who (%)
T9.5 Support for Repack2	Giulia(8%)
T11.2 Maintain CASTOR web site: Keep web site updated	Rosa(10%)
T13.1 Project Management	Sebastien(20%)
T13.2 Tier-1 Support in particular CNAF	Giuseppe(20%)
T13.3 3rd-level support	<i>all</i>
T13.4 Miscellaneous bugfixes	<i>all</i>
T14.1 Castor releases	Giuseppe(7%), Rosa(13%)
T14.1 Meetings and workshops	<i>all</i>



Comments, questions?

