

1 CASTOR related activities 1Q06-3Q06

This note describes the different CASTOR related activities scheduled for the first three quarters of 2006. Where LCG milestones involve CASTOR activities, the latter are constrained by the milestones and will hence take precedence over other tasks, which do not depend explicitly or implicitly (via task interdependence).


2 CASTOR related LCG milestones

The following LCG milestones have a direct impact on CASTOR activities

Milestone	Brief description	Date	Impact
SC3	SC3 service phase	31/12/05	Operation (hardware interventions, software upgrades, reconfigurations) and user support
DRC2	IT internal data challenge. Two internal milestones: 750MB/s CDR and 500MB/s Tier-0 challenge	31/12/05	Hardware setup, configuration and operation. Software problem tracking and policy tuning
SC3 throughput repetition	1GB/s out-rate from CERN to T1. Disk-to-disk transfers	26/01/06	WAN reconfiguration (to LCG network), hardware operation and user support + problem tracking
DRC3	IT internal data challenge. Two internal milestones: 1GB/s CDR challenge, 1GB T0 challenge	30/04/06	Hardware setup, configuration and operation. Software problem tracking and policy tuning
ALICE DC 1GB/s	1GB/s data challenge with ALICE	30/04/06	Re-use DRC3 setup. Hardware operation, problem tracking and user support
SC4a	All required software for baseline services deployed and operational at all Tier-1 and at least 20 Tier-2 sites	28/02/06	CASTOR2 deployed on all T1s.
SC4b	Service Challenge 4 set-up: set-up complete and basic service demonstrated, capable of running experiment-supplied packaged test jobs, data distribution tested.	30/04/06	Hardware setup, configuration and operation. Software problem tracking and user support
SC4-4 start	Service challenge 4: start of stable service phase, including all Teir-1s and 40 Tier-2 sites. The service must be able to support the full computing model of each experiment, including simulation and end-user batch analysis at Tier-2 sites.	31/05/06	Hardware configuration and operation. Software problem tracking and user support.
SC4-4 end	1) 8 Tier-1s and 20 Tier-2s must have demonstrated availability better than 90% of the levels specified in Annex 3 of the WLCG MoU [adjusted for sites that do not provide a 24 hour service] 2) Success rate of standard application test jobs greater than 90% (excluding failures due to the applications environment and non-availability of sites) 3) Performance and throughput tests complete: Performance goal	30/09/06	Hardware configuration and operation. Software problem tracking and user support. All LHC users migrated to CASTOR2

	<p>for each Tier-1 is the nominal data rate that the centre must sustain during LHC operation (see Figure 3): CERN-disk > network > Tier-1-tape.</p> <p>Throughput test goal is to maintain for one week an average throughput of 1.6 GB/s from disk at CERN to tape at the Tier-1 sites. All Tier-1 sites must participate.</p>		
--	--	--	--

The following figure shows the scheduling of the different milestones:



3 Internal LCG-dependent CASTOR milestones

The following internal CASTOR milestones directly depend on the LCG milestones above.

Milestone	Brief description	Date	Dependent LCG milestones
SRM v2.1 delivery	CERN part of the SRM v2.1 development and deployment	31/01/06	SC4a
CASTOR2 client ready for general users	Improvements to the stager_qry command required by all 4 LHC experiments		SC4-4 start
CASTOR2 operation ready for general users	Training of operation team. Automation of installation, configuration and monitoring.		SC4-4 start
LHC users migrated			SC4-4 start

4 Activities

4.1 Overview

The main CASTOR related activities for the coming 9 months have been grouped as follows:

Activity	Brief description
SC3 running	Complete operation and support for the SC3 service phase
DRC2 activities	750MB/s CDR + 500MB/s T0 challenges
DRC3 activities	1GB/s CDR + 1GB/s T0 challenges + ALICE 1GB/s challenge
SRM v2.1 activities	SRM v2.1 development and deployment (CERN part)
SC3 throughput repetition activities	The SC3 throughput repetition (disk-to-disk)

Tape strategy activities	Installation and testing of new robotics and high-end tape drives
Oracle hardware evaluation	Stress testing and evaluation of various Oracle hardware options for the CASTOR name server and the CASTOR2 stager
LHC user migration to CASTOR2	Migration of the general user groups in the 4 LHC experiments
CASTOR2 migration at Tier-1	Migration to CASTOR2 in the Tier-1 institutes
SC4 setup phase activities	Operation and support for the SC4 setup phase
SC4 service phase	Operation and support for the SC4 service phase
Non-LHC user migration	Migration of the non-LHC general user groups
Routine operation	CASTOR1 and CASTOR2 operation and support (2 nd and 3 rd level). Quattorization/lemonization of CASTOR2. Operation training
Strong authentication	Finalization of the developments. Setting up the infrastructure and deployment
Administrative activities	External operation meeting at RAL, Mumbai workshop, CASTOR review
CASTOR2 operation devs	Developments required for routine operation
User migration dev	Developments required for general user migration
Tape operation devs	Developments required for efficient tape operation with the new high-end tape drives
Other developments	Investigation of stager bottlenecks (e.g. LSF tuning) preventing large scale Oracle stress-test. Native support for various protocols (gsiftp, xrootd). RFIO-DPM merge

4.2 Setting up of new CASTOR2 instances

Many of the activities listed in previous section involve the setting up of a new CASTOR2 instances. This is foreseen to be largely automated via Cdb templates. However, castor operation efforts are required for the preparation of the hardware, preparation of Cdb templates, registration of operator alarms and sysadmin procedures. Each CASTOR2 instance also require the setting up of two database instances (one for the catalogue and one for the logging), which require efforts from castor operation for the preparation and configuration of the hardware and the Oracle database administrators to install and configure the database. The effort estimation is:

Who	What	Estimated work
Jan van Eldik, Miguel MARQUES COELHO DOS SANTO	Preparation of hardware, Cdb templates etc	2 Person days
Eric Grancher, Nilo Segura	Configure database	1 Person day


The instances setup for data challenges also requires configuration work for the tape archive (e.g. tape pools, device groups). The additional effort estimation is:

Who	What	Estimated work
Tony Osborne	Configuration of tape pools and device groups etc.	1 Person day

The instances foreseen for next year are:

Instance	Type (DC/prod)	By When
ALICE general users	Prod	10/01/2006
LHCb general users	Prod	19/01/2006
ATLAS general users	Prod	02/02/2006

4.3.8 LHC user migration to CASTOR2




Migration of each of the 4 LHC user communities to CASTOR2. The first migration depends on software developments answering some user requirements expressed by all four experiments (mostly related with the stager_qry). It also depends on the CASTOR user guide to be updated with CASTOR2 information. For each user group the estimated efforts are:

- 3 Person Days (effort drive) for setting up of a new CASTOR instance (see Section 4.2)
- 0.4FTE (fixed duration) for ~20 days of user migration support

The migration is a pre-requisite for the 'SC4-4' start milestone (SC4 service phase).


4.3.9 CASTOR2 migration at Tier-1



A pre-requisite for the SC4 readiness milestone (SC4a) is the SRM v2.1, which only comes with CASTOR2. All CASTOR based Tier-1 institutes must therefore upgrade. The upgrade involves external factors such as hardware and personnel availability at the Tier-1 sites. The required CERN based effort is:

- 1FTE (fixed duration) for 10 days for help with configuration and operational training.

4.3.10 SC4 setup phase activities



Setup and operate a CASTOR2 instance for the SC4 setup phase. The SC4 setup phase ends with the SC4b milestone.

- 3 Person Days (effort driven) for setting up a new CASTOR2 instance (see Section 4.2)
- 0.4FTE (fixed duration) for operation
- 0.6FTE (fixed duration) for software support and problem tracking

Like for the SC3 throughput phase, it is assumed that there is a second phase with point-to-point or other types of tests during May 2005.

4.3.11 SC4 service phase

ID	Task Name	Jun '06				Jul '06				Aug '06				Sep '06				Oct '06				Nov '06	
		28	4	11	18	25	2	9	16	23	30	6	13	20	27	3	10	17	24	1	8	15	22
86	support	Sebastien Ponce[20%]																					

The SC4 service phase ends with milestone 'SC4-4 end', 30/09/2006. It is assumed that it will be operated with the 4 LHC experiments' CASTOR2 instances, for which efforts is already accounted under CASTOR2 operation in a later section. The support load is however expected to be higher the normal CERN user support.

- 0.2FTE (fixed duration) software support and problem tracking

4.3.12 Non-LHC user migration

ID	Task Name	April 2006				May 2006				June 2006				July																
		2	5	8	11	14	17	20	23	26	29	2	5	8	11	14	17	20	23	26	29	1	4	7	10	13	16	19	22	25
88	Develop pla for non-LHC user migration	Tony Osborne[50%],Olof Barring[50%]																												
89	Setting up castor2 stager instance	Jan van Eldik[30%],Miguel MARQUES COELHO DOS SANTO[30%],Eric Grancher[10%],Nilo Segura[10%]																												
90	Non-LHC users migration to CASTOR2	[Bar chart showing migration effort across months]																												

Migration of each of the non-LHC user communities to CASTOR2.

- 10 Person Days (effort driven) for gathering requirements from all remaining CASTOR1 user groups (~20 in total) and develop a plan for the migration.
- 3 Person Days (effort drive) for setting up of a new CASTOR instance (see Section 4.2)
- 0.4FTE (fixed duration) for ~20 days of user migration support

4.3.13 Routine operation

ID	Task Name	2005			Qtr 1, 2006			Qtr 2, 2006			Qtr 3, 2006			Qtr 4, 2006			Qtr 1, 2007			Qtr 2, 2007			Qtr 3, 2007			Qtr 4, 2007			Qtr 1, 2008		
		Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb		
92	CASTOR2 operation training	Olof Barring[40%],Miguel MARQUES COELHO DOS SANTO[20%],Tony Osborne[20%],Jan van Eldik[20%],Veronique Lefebure[20%]																													
93	CASTOR2 updates to user guide	Tony Osborne[70%],Olof Barring[30%]																													
94	CASTOR2 quattorization/lemonization	Miguel MARQUES COELHO DOS SANTO[50%],Jan van Eldik[50%],Alasdair Earl[50%]																													
95	CASTOR1 operation	Tony Osborne[50%],Miguel MARQUES COELHO DOS SANTO[10%]																													
96	CASTOR2 operation	Tony Osborne[20%],Miguel MARQUES COELHO DOS SANTO[20%],Veronique Lefebure[20%]																													
97	CASTOR2 3rd level support	Olof Barring[40%],Giuseppe Lo Presti[40%],Sebastien Ponce[40%]																													
98	COMPASS and NA48 CDR	Tony Osborne[20%]																													

Routine operation encompasses both CASTOR1 and CASTOR2 operation and user support. The CASTOR2 operation includes training of personnel and updates to the user guide. Because there will be at least 5 permanent CASTOR2 instances (one per LHC experiment + one for non-LHC users) plus a couple of instances being setup for the various data and service challenges, it is also important that the automated configuration, installation and monitoring are developed as far as possible.

- 1.2FTE (fixed duration) for operational training (0.4FTE for developing the material and giving the training and 0.2FTE learning effort from each of four the participating 'students').
- 5 Person Days (effort driven) for updating the CASTOR users' guide
- 40 Person Days (effort driven) for quattorization/lemonization the CASTOR2 operation
- 0.6FTE (fixed duration) for CASTOR1 operation. This is expected to ramp down as the number of CASTOR1 instances will decrease when users are migrated to CASTOR2

- 1FTE (fixed duration) for operating CASTOR2. This includes the operational efforts for SC4 service phase (see Section 4.3.11)
- 1.2FTE (fixed duration) 3rd level software support and problem tracking
- 0.2FTE (fixed duration) for operating CASTOR1 for COMPASS and NA48 CDR.

4.3.14 Strong authentication

ID	Task Name	Duration	2006																																																				
			Jan '06							Feb '06							Mar '06							Apr '06							May '06							Jun '06							Jul '06							Aug '06			
101	Develop plan for deployment of strong authentic	20 days	Sebastian Lopienski																																																				
102	Prepare infrastructure for strong authentication	20 days	Sebastian Lopienski																																																				
103	Deployment of strong authentication	40 days	Sebastian Lopienski																																																				

The deployment of strong authentication should take place after all LHC and non-LHC user groups have been migrated to CASTOR2, otherwise there is work required in securing the CASTOR1 stager. As special arrangement may therefore be required for the NA48 and COMPASS CDR, which will probably run with CASTOR1 stagers (see Section 4.3.13).

- 20 Person Days (effort driven) to develop a plan for the deployment of strong authentication
- 20 Person Days (effort driven) for preparing and verifying the infrastructure
- 40 Person Days (effort driven) for deploying the strong authentication

4.3.15 Administrative activities

ID	Task Name	Duration	2006																											
			January 2006							February 2006							March 2006							April 2006						
105	RAL workshop	3 days	Sebastien Ponce, Giuseppe Lo Presti																											
106	Prepare for Mumbai workshop	3 days	Ben Couturier, Sebastien Ponce, Jan van Eldik																											
107	CASTOR review	50 days	Sebastien Ponce[20%], Olof Barring[2]																											
108	Preparation of CASTOR review report	5 days	Sebastien Ponce[50%], Olof B																											

This includes activities not accounted somewhere else like the preparations for two workshops (CASTOR workshop at RAL and the SC4 workshop in Mumbai) and the CASTOR review (as requested by LHCC reviewers in December 2006).

- 2FTE (fixed duration) for three days for the CASTOR external operations workshop at RAL, January 23-25, 2006
- 3FTE (fixed duration) for three days (including two days for travel) for the SC4 Mumbai workshop, February 10-12, 2006
- 0.4FTE (fixed duration) for CASTOR readiness review. It is still unclear what efforts is really required on the CASTOR side so the estimated 0.4FTE is more a placeholder
- 5 Person Days (effort driven) for writing or contributing to the CASTOR readiness review report. Again the real required effort is unknown.

4.3.16 CASTOR2 operation developments

ID	Task Name	Duration	2006																											
			January 2006							February 2006							March 2006							April 2006						
110	Stager leak	10 days	Sebastien Ponce																											
111	disk2disk copy problem	10 days	Sebastien Ponce																											
112	stager qry independent of svcclass	5 days	Sebastien Ponce																											
113	gc tuning	5 days	Sebastien Ponce																											
114	LSF interface improvements	20 days	Sebastien Lopienski																											
115	Automatic DB cleanup at stager restart	5 days	Sebastien Ponce																											
116	Filesystem weight policy by svcclass	10 days	Sebastien Ponce																											

A number of CASTOR2 software improvements required for a smooth operation of the system. Operational workarounds exist for most of the problems but the procedures will not scale well with the number of instances foreseen.

- 10 Person Days (effort driven) for fixing an outstanding memory leak in the CASTOR2 stager. This problem has existed since the beginning of the SC3 service phase. The workaround is to regularly restart the stager (no service interruption since the process is stateless).
- 10 Person Days (effort driven) for fixing a problem with internal file replications between CASTOR2 service classes. A rare but annoying problem for the users because the affected files stay in an intermediate state where they cannot be accessed. The workaround is to regularly check and cleanup the stuck requests from the stager catalogue.
- 5 Person Days (effort driven) for fixing a problem with the stager file-query not correctly supporting disk-pool independent queries. This causes problems for SRM clients querying for disk resident files
- 5 Person Days (effort driven) for tuning the garbage collection policy.
- 20 Person Days (effort driven) for further optimizations in the CASTOR2 – LSF interface and tuning of the LSF configuration.
- 5 Person Days (effort driven) for implementing automatic cleanup of requests and files in inconsistent states after an abrupt shutdown of the stager
- 10 Person Days (effort driven) for implementing filesystem weights per service class.

The above Gantt chart shows the development timelines before leveling of the whole project. Therefore it only shows the estimated effort not the real duration of the developments.

4.3.17 User migration developments

ID	Task Name	Duration	December 2005							January 2006							February 2006							March 2006									
			6	9	12	15	18	21	24	27	30	2	5	8	11	14	17	20	23	26	29	1	4	7	10	13	16	19	22	25	28	3	6
118	stager qry per castor directory	20 days	[Bar]							[Bar]							[Bar]																
119	stager qry returning disk pool occupancy	20 days	[Bar]							[Bar]							[Bar]																
120	MacOS/windows support for CASTOR client	20 days								[Bar]							[Bar]																
121	CASTOR1-2 disk2disk copy	10 days								[Bar]							[Bar]																

A number of CASTOR2 software improvements required for the migration of general users to CASTOR2.

- 20 Person Days (effort driven) for implementing stager_qry per castor directory or regular expressions. The stager_qry currently only supports query by files.
- 20 Person Days (effort driven) for implementing stager_qry for disk pool monitoring (occupancy).
- 20 Person Days (effort driven) for porting CASTOR2 client to MacOSX and Windows
- 10 Person days (effort driven) for developing a procedure to copy files from a CASTOR1 stager to a CASTOR2 stager (instead of recalling them from tape).

The above Gantt chart shows the development timelines before leveling of the whole project. Therefore it only shows the estimated effort not the real duration of the developments.

The two stager_qry developments have been requested by the four LHC experiments. An internal milestone ‘CASTOR2 client ready for general users’ has therefore been defined for their estimated completion date.

4.3.18 Tape operation developments

ID	Task Name	Dec '05		Jan '06			Feb '06			Mar '06			Apr '06			May '06			Jun '06			Jul '06							
		20	27	4	11	18	25	1	8	15	22	29	5	12	19	26	2	9	16	23	30	7	14	21	28	4	11	18	25
123	rmcdaemon/tpdaemon revamp to support large IBM robotics	Ben Couturier[20%],Hugo MONTEIRO CACOTE[80%]																											
124	Immediate tape mark writing	Hugo MONTEIRO CACOTE																											
125	Embedded labels	Ben Couturier[10%],Olof Barring[20%],Hugo MONTEIRO CACOTE																											
126	Measurement of tape position timing map	Charles Curran																											
127	Read/write policies based tape media layout	Olof Barring																											
128	Repack	Felix Ehm																											
129	Recall policies	Olof Barring																											
130	Deployment of new VDQM server	Matthias Braeger[80%],Olof Barring[20%]																											

Various developments for a reliable and efficient use of the new robotics and high-end tape devices.

- 32 Person Days (effort driven) for revamping the rmcdaemon to robustly support large IBM robot. This activity includes the handing over of rmcdaemon/tpdaemon software support to TSI efforts
- 20 Person Days (effort driven) for enable and test immediate tape mark writing in the CASTOR tape position interface.
- 60 Person Days (effort driven) for developing and testing of tape file header and trailer embedded with data file. This is a major development and moves away from the use of tape label standards in CASTOR.
- 20 Person Days (effort driven) for measuring of tape position timing map reflecting the physical serpentine layout of the data on tape
- 10 Person Days (effort driven) for using the tape position timing map (previous point) for optimizing tape migration and recall
- 60 Person Days (effort driven) for developing a new repack application
- 20 Person Days (effort driven) for developing the framework for tape recall policies
- 10 Person Days (effort driven) for deploying the new VDQM server. The deployment foreseen for December 2005 failed and a new deployment has to be scheduled sometime in 2006.

4.3.19 Other developments

ID	Task Name	December 2005				January 2006					February 2006					March 2006																					
		27	30	3	6	9	12	15	18	21	24	27	30	2	5	8	11	14	17	20	23	26	29	1	4	7	10	13	16	19	22	25	28	31	3	6	9
132	Investigation of stager performance problems	Sebastien Ponce[50%],Sebastian Lopienski[50%]																																			
133	Training of new development staff	Sebastien Ponce[50%]																																			
134	xrootd support	Sebastien																																			
135	RFIO-DPM merge	Ben Couturier																																			
136	VOMS integration																																				

This heading groups various important developments that could not be classified elsewhere.

- 20 Person Days (effort driven) for investigating stager performance bottlenecks, improvements of the CASTOR2-LSF interface and LSF configuration tuning. This task is a pre-requisite for the CASTOR2 stager catalogue Oracle stress testing (see Section 4.3.7)

- 1FTE (fixed duration) for ~20 days for training new development staff
- 40 Person Days (effort driven) for adding native support for xrootd disk mover in CASTOR2
- 20 Person Days (effort driven) for merging CASTOR and DPM RFIO implementations
- 20 Person Days (effort driven) for integrating SRM v2.1 with VOMS. A much large task is the full integration of all CASTOR2 interfaces with VOMS. For the moment there is no estimate of that effort.

5 All activities and milestones together

The full plan for all groups of activities and milestones listed in previous sections but without efforts leveling is shown below.

