



Planning and Resources

Castor Development Team

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

Castor Readiness Review – June 2006



Outline



- ❖ Current status of CASTOR-2 software
- ❖ Evolution of development resources since 1-2005
- ❖ Summary of development plans for 2006



Outline



❖ Current status of CASTOR-2 software

- Stager, scheduler
- Clients, supported protocols
- Name server
- Distributed Logging Facility
- Tape Area
- SRM
- Security

❖ Evolution of development resources since 1-2005

❖ Summary of development plans for 2006



Current status of Castor-2 SW (1)



First Castor-2 release:

- “Released” Q1 2005
- Used for Alice DAQ tests in March 2005

❖ Current software release: Castor-2.1.0-0

- Released May 15, 2006
- Deployment at CERN, CNAF (RAL, PIC, ASGC to follow)
- Fully site-independent (customization via cfg files)
- Integrated into automated systems mgmt + monitoring (Quattor, Lemon)
- Mature architecture/design: focus now is to address specific bug's/RFE's

❖ Platforms:

- Support for SL3/4 Linux (i386 client&server; x86_64&ia64 client)
- Packaged using native RPM format
- no MacOSX, nor other UNIXes
- Windows port of client software started (voluntary external contribution without commitment)



Current status of Castor-2 SW (2)



❖ Stager, Scheduler

➤ Stager Database

- Based on Oracle 10g (10.1/10.2)
 - Scalability is not a major concern, but HW reliability is
 - Very good support from IT/DES
- PostgreSQL support frozen; OracleXE tried out

➤ Configurable policies

- For recall, migration, GC, I/O scheduling
- Support for different storage classes
 - Volatile (GC, no migration), Durable (no GC, no migration) and Permanent (GC, migration)

➤ Advanced migration/recall features

- prepareToGet/Put : session management
- "getNext" : streaming mode

➤ External scheduler

- Only LSF supported right now
 - LSF plug-in scheduler provided
 - Bottleneck limiting to ~ 10 file accesses/s
- MAUI support dropped



Current status of Castor-2 SW (3)



❖ Protocols and Clients

➤ Protocols:

- RFIO and ROOTD fully integrated in CASTOR-2 as 'internal' protocols
- Support for XROOTD recently added (collaboration with A. Hanushevsky)
- GridFTP v1 supported but not as 'internal' protocol (uses RFIO internally)

➤ RFIO issues:

- still different in DPM and CASTOR-2
- IT/GD has started to work on a common framework

➤ Clients:

- Full set of command line clients exist for all CASTOR-2 services
- Stager, name server, tape, protocols
- Packaged as independent RPM sets
- Can coexist with CASTOR-1 commands/RPM's
 - Backwards compatible, except stager commands



Current status of Castor-2 SW (4)



❖ Name Server

- Using DPM(/LFC) name server code base
 - Collaboration with J-P Baud (IT/GD)
- DPM name server is itself an evolution of the old CASTOR-1 name server
 - Provides POSIX ACL support (needed for SRM2.1)
 - Many other new features e.g. strong authentication, symlinks, checksums, transactions/sessions
- Re-added CASTOR-specific extensions (e.g. tape/segment related)
 - ... no common DPM/CASTOR-2 CVS repository yet
- Support for MySQL and Oracle



Current status of Castor-2 SW (5)



❖ DLF (Distributed Logging Facility)

- Major redevelopment since 11/05 for improving performance and reliability
- Oracle 10g, MySQL
- Used by all CASTOR-2 server-side components
- Web GUI available
 - Browse DLF logs
 - Browse stager DB contents
- Latest release additions, not yet deployed:
 - Partitioning (Oracle)
 - Archiving (Oracle)
 - Performance optimisations



Current status of Castor-2 SW (6)



❖ Tape Area

- RTCOPY, tpdaemon, VMGR
 - coming from CASTOR-1
 - RTCOPY protocol enhanced for dynamically add files to running requests
 - Support for new-generation tape drives (LTO3, T10K, 3592B)
- VDQM
 - rewritten but not yet deployed
- Repack
 - Rewritten from scratch for CASTOR-2; development completed and waiting for rollout



Current status of Castor-2 SW (7)



❖ SRM interface

➤ SRM-1:

- Developed for and using during SC3; used unmodified during SC4
- Implemented permanent/durable storage selection via TURL (for RFIO and ROOT), as agreed in Mumbai

➤ SRM-2:

- SRM-2.1 interface developed in collaboration with RAL
- Started 8/05, full version available since 2/06
- Supporting data transfers, pinning, relative paths, directory functions, global space reservation, VOMS integration, etc
- Uses new nameserver ACL functionality
- LCG decision to use SRM-2.X only after SC4
- SRM-2.2 specification underway (FNAL workshop)



Current status of Castor-2 SW (8)



❖ Security

- Per-file ACL's at namespace level (name server)
- Restricted access to disk servers for internal protocols
- Strong authentication under work
 - Castor Security library available (now maintained by IT/GD), plugin modules for KRB5, GSI
 - strong authentication to be used in client-server communication (RFIO, name server, stager)
 - Deployment planned for Q4 2006, after performance evaluations
- No privacy protection (data encryption) planned
- No internal support for VOMS



Outline



- ❖ Current status of CASTOR-2 software
- ❖ Evolution of development resources since 1-2005
- ❖ Summary of development plans for 2006



Evolution of Resources since 1/05



Developer	From	To	%	Main development responsibilities
Olof Barring	1999	now	100->10	Castor-1, SRM, tape area maintenance
Hugo Cacote	11/05	now	60	Tape driver software
Felix Ehm (tech student)	11/05	07/06	100	Repack
Rosa Garcia Rioja	03/06	now	100	GridFTP/xrootd, 64-bit port, web pages
Giuseppe Lo Presti (INFN fellow)	03/05	now	100	SRM-2, stager + DB code, ETICS integration
Sebastian Lopienski	11/05	now	50	Strong authentication
Sebastien Ponce	09/03	now	100	Project leader+main architect, stager, scheduler
Giulia Taurelli	03/06	now	100	Test suite, RFIO, stager
Dennis Waldron	11/05	now	50	DLF, expert system
Matthias Braeger (tech student)	03/05	12/05	100	VDQM re-engineering
Ben Couturier	01/02	02/06	50-100	Tape HW, name server, SRM-1, client part
Jean-Damien Durand	1999	05/06	100	Castor-1, stager, LSF plugin, expert system, monitoring
Emil Knezo (fellow)	07/02	06/05	100	gridFTP/Castor integration, SRM
Victor Kotlyar, Vitali Motyakov, Victor Zhitsov (IHEP+JINR)	07/03	now	1 -> 0.5 FTE	DLF, DLF GUI, stager, DB GUI, Windows client
Jiri Mencak, Tara Shah, Shaun de Witt (RAL)	04/05	now	2 -> 0.5 FTE	SRM-2



Evolution of Resources since 1/05



Developer	From	To	%	Main development responsibilities
Olof Barring	1999	now	100->10	Castor-1, SRM, tape area maintenance
Hugo Cacote	11/05			
Felix Ehm (tech student)	11/05			
Rosa Garcia Rioja	03/06			
Giuseppe Lo Presti (INFN fellow)	03/05			
Sebastian Lopienski	11/05			
Sebastien Ponce	09/03			
Giulia Taurelli	03/06			
Dennis Waldron	11/05			
Matthias Braeger (tech student)	03/05			
Ben Couturier	01/02			
Jean-Damien Durand	1999			
Emil Knezo (fellow)	07/02			
Victor Kotlyar, Vitali Motyakov, Victor Zhitsov (IHEP+JINR)	07/03	now	1 -> 0.5 FTE	DLF, DLF GUI, stager, DB GUI, Windows client
Jiri Mencak, Tara Shah, Shaun de Witt (RAL)	04/05	now	2 -> 0.5 FTE	SRM-2

- ❖ High departure rate in 2005/6, leaving CASTOR-dev with two core members
- ❖ New arrivals in Q1/Q2 2006
- ❖ Reallocation of section members to CASTOR
- ❖ Newcomers to be trained by existing team members
 - but weak training material!
 - Overlapping with first wider production rollouts of CASTOR-2
- ❖ Consequences..:
 - Delays of tasks and milestones
 - Drop non-vital activities and concentrate on core developments

- "senior" team member
- "junior" team member (<=6 months)
- former team member
- ext. collaborator




Outline



- ❖ Current status of CASTOR-2 software
- ❖ Evolution of development resources since 1-2005
- ❖ Summary of development plans for 2006



Summary of dvp plans for 2006



■ medium, ■ high, ■ low importance

❖ The complete and up-to-date planning can be found on the [castor web](#)



Summary of dvp plans (2)



- TG1, TG2 (known bugs and testing)

Task	Start;duration progress	Who (%)
T1.1 Memory leaks in stager; currently worked around via crond restarts	15/7/06; 4w	Sebastien(10%)
T1.2 Bad database states due to wrong PrepareToPut/get sequence handling	10/6/06; 10d	Sebastien(10%)
T2.1 Test recovery after DB restart: Daemons to detect if stager DB connection was lost	15/5/06; 3w 30%	Giulia(100%)
T2.2 Test recovery after DLF restart: Daemons to detect if connection to DLF was lost	1/8/06; 4w	Dennis(50%)
T2.3 Test DB hardware configurations for stager (including RAC and DataGuard)	1/10/06; 4w	?
T2.4 New DLF testing and deployment (improved performance, resilience, support for archiving)	1/7/06; 4w	Dennis(50%)

□ medium, □ high, □ low importance



Summary of dvp plans (3)



- TG3, TG4 (Stager, Scheduler)

Task	Start;duration progress	Who (%)
T3.1 Signal handling on client part (gracefully handle signals like SIGINT)	1/1/07; 3w	?
T3.2 stager_query improvements: reduce load on stager by minimising query overhead (e.g. due to regexps)	1/9/06; 2w	Sebastien(10%)
T3.3 Automatic cleanup at restart: remove stale DB entries after crashes of components	1/10/06; 4w	?
T3.4 SvcClass specific selection policy: allow different filesystem selection policies per service classes and disk pools	1/10/06; 3w	?
T3.5 Recall policies: allow for a smaller scheduling and grouping of recall operations	1/10/06; 3w	?
T3.6 Filesystem GC trigger optimisation	10/5/06; 2d 80%	Giuseppe(100%)
T4.1 LSF limitations: study and improve limitations in LSF plugin for higher job/s throughput rates	1/7/06; 4w	Sebastien(10%)

□ medium, □ high, □ low importance



Summary of dvp plans (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	15/6/07; 3m	Giulia(20%)
T5.2 GridFTP: add support for GridFTP2 as internal CASTOR-2 protocol	1/8/06; 3m	Rosa(90%)
T5.3 XROOTD: assist xrootd developers in integrating/testing support for XROOTD as internal CASTOR-2 protocol	1/4/06; 3m 50%	Rosa(40%)
T6.1 Strong authentication: Finish the implementation of strong authentication in CASTOR-2; define a rollout plan	1/3/06; 6m 15%	SebastianL.(40%)
T6.2 Accounting: Implement accounting in CASTOR-2 and accounting-based access policies	1/4/07; 2m	?

□ medium, □ high, □ low importance



Summary of dvp plans (5)



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

Task	Start;duration progress	Who (%)
T7.1 DLF improvements: improve DLF performance and client resilience; add database archiving	1/2/06; 4m 95%	Dennis(50%)
T7.2 DLF archiving GUI: enhance DLF GUI for managing DLF archives	1/8/06; 1m	Viktor K.(100%)
T7.3 SRM-2 updates, deployment: modify existing SRM-2 implementation as agreed by WLCG for post-SC4	1/5/06; 5m 20%	Shaun(30%)
T7.4 New VDQM: finish development, test and deploy new VDQM	15/6/06; 3m	Giulia(80%)
T8.1 CLIPS to Perl transition: replace complex CLIPS-based expert system by a simpler perl script	1/6/06; 3w	Dennis(50%)
T8.2 rmmaster: Redesign rmmaster and split functionality into separate and stateless monitoring and scheduler interface parts	1/1/07; 3m	?
T8.1 CUPV: redesign old/unmaintained user privilege validation module, reduce complexity and simplify authorization setup	1/4/07; 2m	?

□ medium, □ high, □ low importance



Summary of dvp plans (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	1/2/06; 6m (90%)	Felix(100%)
T9.2 Makefiles and packaging: Provide a clean, modular, well-organised build/packaging schema	15/6/06; 3m	Giuseppe(20%)
T9.3 Automation of releases and tests: automatic regular builds and execution of test suites using ETICS framework	15/7/06; 2m	Giuseppe(20%)
T10.1 Port of servers to 64bits: port CASTOR-2 server parts to 64bit (stager, scheduler, SRM, etc)	1/6/06; 6w	Rosa(50%)

□ medium, □ high, □ low importance



Summary of dvp plans (7)



- TG11 (Documentation, publications)

Task	Start;duration progress	Who (%)
T11.1 New CASTOR web site: redesign CASTOR web site following ELFms look-and-feel, update outdated information	1/5/06; 1m 60%	Rosa(50%)
T11.2 Maintain CASTOR web site: Keep web site updated	1/6/06; 6m	Rosa(10%)
T11.3 Guides: Provide missing CASTOR-2 guides (User, Admin, Developer)	1/4/07; 2m	?
T11.4 Articles: publish articles on CASTOR-2 (architecture, deployment, HA, ..)	1/7/06; 1m	Giuseppe(50%)

□ medium, □ high, □ low importance



Comments, questions?

