



CASTOR SRM CCRC-08

Shaun de Witt
CCRC08 Post Mortem
12 June 2008



Introduction

- Problems in CCRC-08
 - And what we are doing about them
- Positives
- Future Developments
- Release Procedures



Type of Problem

- Database problems
- Network problems
- CASTOR interaction problems
- Pure SRM problems



Problems - Database

- Deadlocks
 - Daemon/Daemon deadlocks
 - Fixed in next release
 - Some code moved to database
 - Server/Daemon deadlocks
 - Solution identified and part implemented
- Too many connections
 - Configuration/resource issue



Problems - Network

- Intermittent CGSI errors
- Terminal CGSI errors
 - Reasons unclear
 - May be related to deadlocks
 - Using older version of VDT and cgsi-plugin
 - Configuration parameters may help
- SRM 'lock-ups'
 - Debugging indicates all threads stuck in recv()
 - Send and Receive timeouts added in latest release



Problems - Interface

- Interactions with CASTOR
 - Behaviour when CASTOR is slow
 - Resilience added by using TIMEOUT parameters on stager call
- Bulk Removal requests
 - CGI-gSOAP errors in bulk srmRm
 - But does work
 - Load related
- Pinning
 - Default pinning too aggressive
 - Plus some miscommunication problems resulted to too heavy a pin being applied
 - Fixed
- Logging
 - Log files contain most information
 - But too distributed
 - Difficult to trace requests between front and back ends
 - 90% solution implemented – last 10% close.



Problems - SRM

- Occasional core dumps from back end
 - Causes identified and corrected



Positives

- Request rate
 - At RAL on 1 cms front end with 50 threads:
 - 21K requests/hr
 - Distribution of type of request not known.
- Processing speed
 - Again using CMS at RAL
 - Put/Get requests in 1-5 seconds
 - Does not include authentication overhead
- Front end was quite stable



Future Developments

- Move to SL4
 - URGENT
- Move to castor clients 2.1.7
 - URGENT
 - New number is major.client.release
- New MoU
 - URGENT



Release Procedures

- Following problems just after CCRC
 - Srm seemed to pass all tests
 - But daemon failed immediately in production (CERN and RAL)
- Brought about by a 'simple' change which only affected recalls when no space token was passed.
 - Clear need for additional tests before release
 - Public s2 not enough



Pre-Release Procedures

- (Re) Developing shell test tool which will be delivered with the SRM.
 - To include basic tests of all SRM functions
 - Will include testing of tape recalls if possible (i.e. not if only using a Disk1Tape0 system)
 - New tests can be added when we find missing cases.
 - Used for regression testing
- Looking at running FULL s2 test suite
- Both before running on PPS



Pre-Release Procedures

- As now, s2 tests will be run over 1 week to try and ensure stability
- PPS made to look more like production set up
- Only after this will the release be made available for use.
- New procedures will be used for 2.7 releases.