



Storage-ware Review: Problems Encountered & Roadmap

F. Donno

CERN

04 March 2008

Sources



- CCRC08 e-log
 - <https://prod-grid-logger.cern.ch/elog/CCRC'08+Logbook/>
- GGUS: <http://www.ggus.org>
- srm-deployment@dcache.org
- dCache request tracker: <http://www.dcache.org/rt/>
- hep-project-castor-srm@cern.ch
- CASTOR request tracker:
<https://savannah.cern.ch/bugs/?group=castor>
- S2 test results and e-mails

CASTOR



- Invalid checksum value returned by the CASTOR gridftp2 server (reported by CMS on 05/02)
 - **FIXED in 1.13-11 (07/02)**
- Gsiftp TURLs returned by CASTOR are relative (reported by S2 and CMS on 06/02)
 - gsiftp://host:port/path/file relative
 - gsiftp://host:port//path/file absolute
 - **FIXED in 1.13-11 (07/02)**
- Unable to map request to space for policy TRANSFER_WAN (reported by CMS on 07/02)
 - **FIXED in 1.13-13 (08/02)**

CASTOR



- The srmDaemon attempts to free an unallocated pointer and crashes (reported by CNAF)
 - **FIXED in 1.13-14 (14/02)**
- Some of the database at CERN have shown an index to be missing (found by S2).
 - **FIXED in 1.3.10-1 (15/02)**

CASTOR



- Insufficient user privileges to make a request of type StagePutDoneRequest in service class 'atldata' (reported by S2 and ATLAS on 19/02)
 - ☺ PutDone executed by and allowed for (root,root) **To be fixed**
 - ✓ **Workaround provided on 23/02**
- Missing access control on spaces based on voms groups and roles (reported by ATLAS/LHCb on 19/02).
 - ✓ **Followed by Storage Solution WG**
- Could not get user information: VOMS credential ops does not match grid mapping dteam (reported by S2 and CNAF on 21/02)
 - ☹ **Not yet understood**

CASTOR



- Error creating statement, Oracle code: 12154 ORA-12154: TNS:could not resolve the connect identifier specified (reported by S2 and CNAF on 12/02)
 - **Not yet understood**
 - ☞ It happens at service startup. A restart cures the problem
- Server unresponsive at RAL? - Space token ATLASDATADISK does not exist (reported by S2 and ATLAS on 28/02)
 - ✓ **Number of threads increased from 100 to 150 (28/2)**

CASTOR



- 10 software problems reported
- 6 problems fixed (in 2-3 days average)
- The problems only affected a small part of the service:
 - srmCopy transfers from FNAL, CMS transfers to CERN, configuration of Atlas MSS + SRM at CERN
- Only 1 important outstanding issue followed by SSWG
- Operational issues are not listed (see talk from HR):
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/CastorOperationsLogbook>
- Both developers and operation people very responsive.

- SRM 1.3-15 for SL4 in test now:
<https://twiki.cern.ch/twiki/bin/view/FIOgroup/RoadMap>
- CASTOR 2.1.7 will be released the first week of April.

dCache



- srmcopy version 1 to version 2 in push mode does not specify TargetStorageType=PERMANENT (reported by ATLAS on 10/01)
 - **FIXED in 1.8.0-12 (22/01)**
- globus-url-copy fail if implicit space reservation had been enabled for non SRM transfers (reported by IN2P3 and FZK sites 22/01)
 - **Workaround: defining specific links outside a linkgroup for non SRM transfers.**
- The SRM reserved space calculation is sometime incorrect (reported by FZK and others on 22/01)
- FTS transfers fail for long transfer (reported by NDGF on 15/01)
 - **FIXED in 1.8.0-12p2 (23/01) and 1.8.0-12p4 (30/01)**

dCache



- Passing the token to the dCache migration script (reported by IN2P3 on 20/12)
 - **FIXED in 1.8.0-12p5 (12/02)**
- Numerical vs. string tokens in srmCopy (reported by CMS and ATLAS on 01/02)
 - **FIXED in 1.8.0-12p5 (12/02)**
- dCache retries SRM 1.1 GET and PUT requests if the protocol is not supported (reported by SARA and PIC on 10/02).
 - **FIXED in 1.8.0-12p5 (12/02)**

dCache



- Failed to create, got error return code from pnfs: File exists (found by S2 and reported by ATLAS and LHCb on 06/02)
 - **FIXED in 1.8.0-12p6 (22/02)**
- Request to increase space sizes at the administrative level (reported by Edinburgh/FZK on 12/02)
 - **FIXED in 1.8.0-12p6 (22/02)**
- SrmCopy push transfers from dCache to Castor failing due to file size mismatch (reported by CMS on 14/02)
 - **FIXED in 1.8.0-12p6 (22/02)**
- gPlazma seems to mix DN and Role if presented at the same time by different requests (reported by LHCb on 10/02)
 - **FIXED in 1.8.0-12p6 (22/02)**

dCache



- Channel closed, may be due to thread interrupt (reported by IN2P3, NDGF and S2 on 12/02)
 - It happens when Berkley DB used on the pool
 - **No fix available yet. Workaround: restart the pool**
 - **Fix will be available in 1.8.0-13**

- Files deleted from PNFS but not yet migrated to tape are not deleted from the pool. Files are precious by default. (reported by LHCb on 21/02 and FZK and IN2P3 on 27/02)
 - **FIX available:** enable the switch to remove files from the pool even when they are precious. Files are removed only when found in the pnfs trash.

dCache



- Failed to create a control line, Failed open file in the dCache. (reported by LHCb on 23/02)
 - It happens when SRM server and gsidcap server run on the same machine. This is not a recommended setup.
- gsidcap connections closed (reported by LHCb on 20/02)
 - Timeout (2hours) set on the duration of the gsidcap connection.

dCache



- Spaces unprotected from generic users usage (reported by ATLAS/LHCb on 25/01)
- Spaces can be released by generic users (reported by CMS on 28/01)
 - **This is a new addition to the WLCG SRM v2.2 MoU Followed up within the SSWG**
- Selecting spaces with PrepareToGet or BringOnline
 - **This is a new addition to the WLCG SRM v2.2 MoU**
 - ✓ **Followed up within the SSWG**

dCache



- gsidcapd and xrootd can not write into spaces managed by the SpaceManager (reported by IN2P3 on 12/2007 and Alice on 22/02): Server error message for [2]: "No write pools configured for ..." (errno 19).
 - **Fix not yet available**
 - **Workaround: defining specific links outside a linkgroup for non SRM transfers**
 - ✓ gsidcap and spaces will be fixed soon (but not yet in 1.8.0-13)
 - ✓ xroot can very likely only be fixed if xroot is authenticated, otherwise it is not possible to protect tokens since the info about VO or the FQAN is missing

dCache



- ~20 software problems reported
- almost all fixed (very good developers response – patch applied to specific sites in production [FNAL])
- The problems affected transfers, data access, cleaning up of spaces
- 2 new important features followed by SSWG
- Operational issues are not listed with some exception (developers have followed closely the sites through weekly phonecalls)
- Site admins have been also very responsive and collaborative.

dCache



- dCache 1.8.0-12 is a special CVS branch for CCRC08. Developers would like now to maintain only main trunk.
- dCache 1.8.0-13 (main trunk) in development. Releasing date not yet known.
- For a list of fixes outside of the CCRC08 branch, please look here:
 - <http://trac.dcache.org/trac.cgi/report/10>
- The dCache team does not force sites to upgrade. The decision is left to sites and experiments.

DPM



- Default ACLs on directories do not work (reported by ATLAS on 13/02)
 - **FIXED in 1.6.7-4 (certified)**
- Slow file removal (reported by ATLAS on 22/02):
 - **ext3 filesystems much slower than xfs for delete operations (2048 files of 1.5GB removed in 90minutes against 5 seconds of xfs – tests performed on the 25/02)**
- DPM 1.6.10 is being certified and will probably be the release available for CCRC08 in May.
 - **The list of fixes/features can be found here:**
<https://savannah.cern.ch/patch/index.php?1605>
 - support for MacOSX
 - group writable directories when SRM started with umask 0
 - IPv6 support

StoRM



- Wrong gsiftp TURLs returned for srmCopy operations (reported by S2 and ATLAS on 12/12)
 - **FIXED on 22/01, available in 1.3.19-1 (18/02)**
- GENERAL_FAILURE - SRM protocol violation! Cannot do an srmPtG of a SURL that is not local! (reported by ATLAS on 13/02)
 - **FIXED on 13/02 available in 1.3.19-1 (18/02)**
- srmPrepareToGet returning wrong size for files \geq 2GB (reported by ATLAS on 16/02)
 - **FIXED in 1.3.19-1 (18/02)**
- New release 1.4.0 in the first week of April 2008
http://storm.forge.cnaf.infn.it/doku.php?id=documentation:storm_release_plan

Conclusions



- Very positive and constructive experience. Problems were found but the response was good. Common feeling that the storage services were functional. However, the storage instances were not stressed and not all features tested.
- Some missing functionality identified: see Storage Solutions working group
- Many sources of information to retrieve and understand the history of a problem. Hard to follow in detail.
 - GGUS tickets are still open (response time ?)
 - e-log threads not “closed” (history ?)
 - a lot in e-mail and not tracked
- Problems generally reported to many people at the same time. Sometimes still did not reach the right target! Sometimes hard to debug.