


# GSSD - report from pre-GDB

GDB - CERN  
10 October 2007

Flavia Donno  
IT/GD, CERN





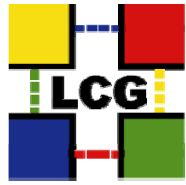
## The meeting

- Good attendance:
  - Experiment representatives
  - Sites
  - Developers
  - OSG participation
    - Promising results of read/write tests using dCache 1.8 and SRM v1.1
    - Participation in Glue Schema discussions
    - Validation of high-level tools




## ATLAS SRM v2.2 testing experience

- Tests performed without modifications to ATLAS DDM
- **No space token used, no pre-staging**
- Export of small (~100KB) and large files (~3GB) from CASTOR@CERN SRM v1 to dCache@NDGF and StoRM@CNAF
- Reached 75% efficiency at NDGF and 95% at CNAF.
- The cause of problems were **staging timeouts at CERN**.
- Some problem observed with StoRM when **operating close to the quota**. The problem is understood and fixed.
- In the process of adding IN2P3, FZK and BNL
- Preparing to test high-level tools/APIs (gfal/lcg-utils)
- Plan to test until it is allowed



## Status of Tier-1 SRM v2.2 test sites

- Static information can be found here:
  - <https://twiki.cern.ch/twiki/bin/view/LCG/GSSDSitesStatus>
- Most sites are OK. **CNAF and RAL CASTOR** instances to come
- Tests performed are being **integrated with new site monitoring**
- Working on **dynamic view** through SAM tests
- New releases foreseen for dCache and CASTOR/SRM. Coordination for upgrades to avoid disruption of ATLAS/LHCb tests



# Schedule for SRM v2 deployment: production

- <http://trac.dcache.org/trac.cgi/wiki/dcache18Deployment>

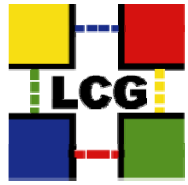
## Tier I's

The table below summarizes milestones in the deployment process of dCache 1.8 (SRM2.2) for the Tier I sites. Please note that at any time, any site may change the scheduling without further notice to dCache.org.

Site	Week of Upgrade	Comment
NDGF	Oct 29	Agreed
gridKa	Nov 5	Agreed. Downtime will be Nov 6, Tuning Nov 7
nothing yet	Nov 19	-
IN2P3	Nov 26	Agreed
RAL	Dec 3	Agreed
BNL	Dec 10	rather sure
PIC	Dec 17	Agreed


**Fermilab** is waiting for a specific dCache feature which we expect to be available mid of Decemeber. So Fermilab will upgrade either end of December or beginning of January.

**Triumf** will decide end of this week (Oct 5).



## Schedule for SRM v2 deployment: production

- CASTOR
  - CERN : end of October 2007
  - RAL : end of November 2007
  - CNAF : as soon as problems with test instance are understood
  - ASGC ?



## Tier-2s

- dCache

### Tier II's

Theoretically, Tier II's may upgrade as soon as the dCache 1.8 production version will be available, which will be the second week of November. Nevertheless we would recommend to watch the upgrades of the Tier I's carefully to estimate the efforts for your site. Furthermore, we would recommend to consider to participate the Edinburgh workshop in November to learn more on how to setup dCache according to the requirements of the experiments.

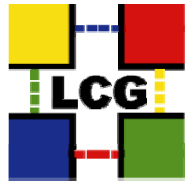
**OSG Tier II's** : End of February, dCache 1.8 will be packaged with the OSG VDT distribution, which is the point in time from which on OSG will officially support dCache 1.8.

- DPM

- Tier-2 can start upgrading as of now. Please, check the workshop in Edinburgh

- StoRM

- ??



## Tier-1/Tier-2 Workshop in Edinburgh


- Organized by GridPP/WLCG-GSSD, hosted by NeSC, **13-14 November 2007**
- Goal: **to prepare Tier-1s and Tier-2s for SRM v2.2 deployment in production**. SRM concepts, Glue schema, site validation procedures, monitoring, tutorials, and hands-on
- Participation by **invitation only** (all Tier-1s, Tier-2s selected by the experiments)
- Main focus on **dCache and DPM**





## Glue Schema discussions

- **Many discussions** on GSSD mailing list about the object classes and attributes exposed by the current Glue 1.3 (Storage Areas, Spaces, Free/Used space, etc.)
- **GSSD focuses on deployment issues.** Working group to start discussions on Glue schema v2.0 for storage services coordinated by Laurence Field
- Within 2 weeks the **example made available** in the GSSD pages will be finalized with only information used by the current clients. This will be the reference for the static information providers created by YAIM, new GIPs, and the SAM validation tests for SRM v2.2
- This example and the new publication and validation tools **will be described during the Tier-1/Tier-2 workshop in Edinburgh**



## Status of the SRM v2 implementations


- **CASTOR:**
  - behavior of PrepareToGet when token is not provided. Agreement reached within CASTOR team. No coding started.
  - Copy has been implemented in current release. S2 tests not successful. Investigating
  - ChangeSpaceForFiles/PurgeFromSpace: design agreed. Ready ASAP.
  - Next release October 15<sup>th</sup>, 2007
  - Working on support model for external sites
- **dCache:**
  - New upgrade procedure from 1.7 to 1.8: no conversion needed for control files
  - Responsibilities for software development agreed at management level
  - Well in line with internal and WLCG schedule
  - Introduced databases on pools instead of control files
  - New release 1.8-19: high frequency of releases from now till production version available
- **StoRM:**
  - Current release 1.3.16: installable via YAIM and APT (rpm and Quattor)
  - New version 1.3.16-2 will be available soon. It fixes minor issues
  - Version 1.4 will be available end of November: srmCopy fully implemented, support for regular expressions for FOAN, fixes for remaining issues reported by S2
  - Independence of spaces and paths will be available after 1.4



# dCache SRM v1 performance

## CMS PhEDEx - Transfer Rate


132 Hours from 2007-10-03 14:00 to 2007-10-09 02:00 UTC



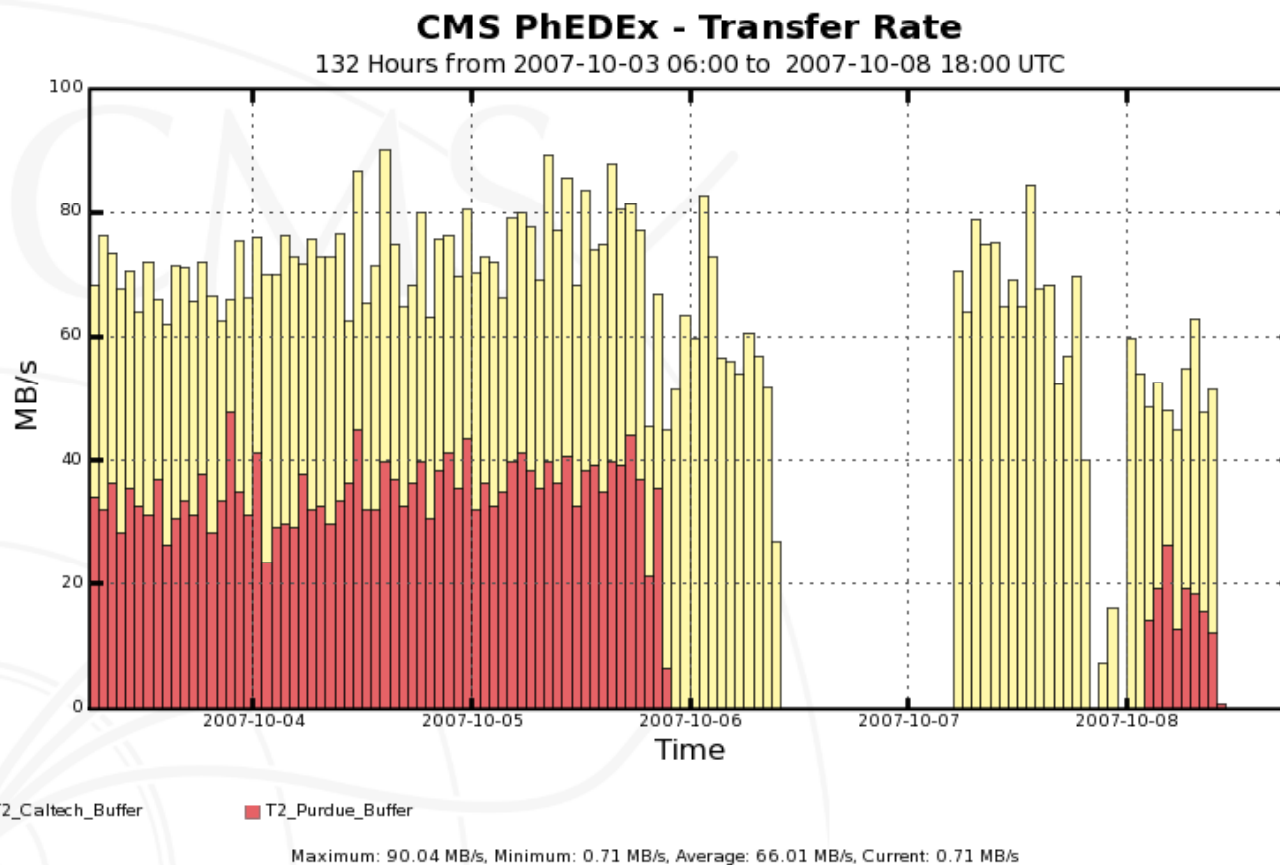
Maximum: 69.69 MB/s, Minimum: 0.71 MB/s, Average: 27.68 MB/s, Current: 1.42 MB/s

Tests performed by CMS at UCSD

Read tests over 132 hours




## dCache SRM v1 performance



Tests performed by CMS at UCSD

Write tests over 132 hours

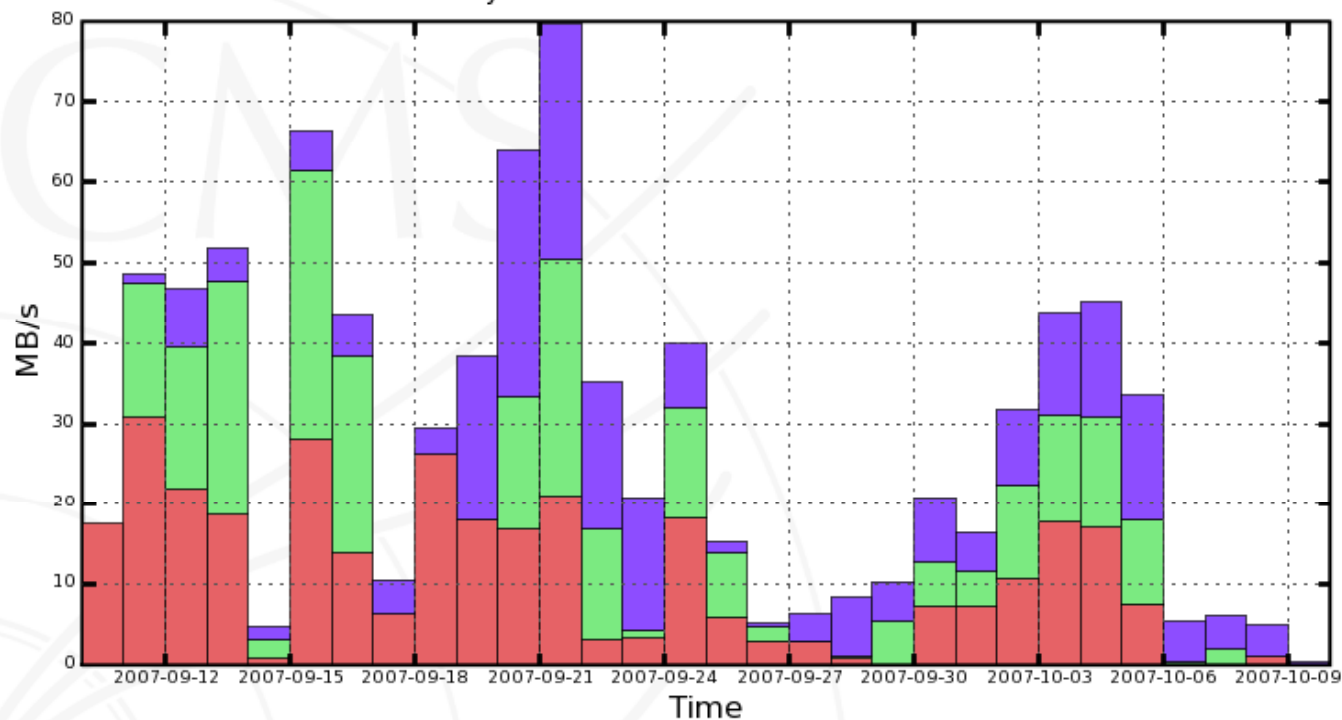
There's a hole. This is fall-out from switching to the bi-directional adapter. Not clear exactly what happened.



# dCache SRM v1 performance

## CMS PhEDEx - Transfer Rate

30 Days from 2007-09-10 to 2007-10-10 UTC



T1\_FNAL\_Buffer

T2\_Caltech\_Buffer


T2\_Estonia\_Buffer

T2\_Purdue\_Buffer

Maximum: 79.61 MB/s, Minimum: 0.24 MB/s, Average: 28.37 MB/s, Current: 0.24 MB/s


Tests performed by CMS at UCSD

Read tests over 30 days



# dCache SRM v1 performance

**CMS PhEDEx - Transfer Rate**  
30 Days from 2007-09-09 to 2007-10-09 UTC



Tests performed by CMS at UCSD

Write tests over 30 days


T2\_Caltech\_Buffer

T2\_Estonia\_Buffer

T2\_Purdue\_Buffer

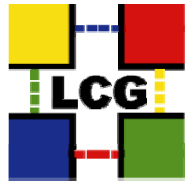
Maximum: 73.49 MB/s, Minimum: 1.75 MB/s, Average: 35.99 MB/s, Current: 21.87 MB/s

We do see a fair number of orphans on the teststand. We are still unclear at this point if this is comparable in scale with what we see in v1.7 of dcache, or if it's worse.



## Status of gfal/lcg-utils


- **Main features** requested by the experiments **available** in the new releases gfal 1.10.x and lcg-utils 1.6.x.
- Pre-stage, pin and release available in gfal and its python library
- No BDII dependency
- **Proven to work in OSG**
- **Plans:**
  - Short term: return version number
  - Medium term: get VO from proxy, rename functions, thread-safe versions



## Experiment requirements


- Availability of **checksum**
- Expose **request ID** through gfal/lcg-utils
- **ReleaseFiles** without specifying request ID
- **Srm-\* like commands** available also in OSG environment (lcg-utils "light" should do the job - possibility to install a tar in experiment area on OSG nodes)
- **Strategy to check consistency** between the content of the **SE** and content of the **site catalogues**
- Possibility to check what has happened in a **time interval**
- **What else ?** Experiments are invited to provide GSSD with the list of requirements as soon as possible






## Use of namespace in transition period (SRMv1 to SRMv2)

- dCache sites used different paths for different storage classes with SRM v1
- Cannot move to new schema with paths independent of spaces before all Tier-1s have migrated to SRM v2
- Sites publish paths in the information system but ...
- ... experiments do not seem to be dependent on that
- **Strategy: use SRM v1 paths till ALL Tier-1s have migrated to SRM v2**



## Tier-2 space organization for experiments in production

- **Input received so far only by CMS**
- More experience with spaces and with SRM v2 at sites needed in order to come up with a strategy
- CMS case:
  - About 200TB at an average Tier-2
  - Used for both production and local users activities
  - It is up to the site to decide if space should be split or not
  - In case, at least 2 space tokens needed: CMS\_PROD and CMS\_USERS
- **We will learn more during the workshop in Edinburgh** where CMS/ATLAS Tier-2s have been invited to present their setup and the experiments requirements for their sites.



## Next GSSD F2F meeting

- November 6<sup>th</sup> 2007
  - Reporting experience during SRM v2 deployment in production
- Tier-1/Tier-2 workshop in Edinburgh, 13-14 November 2007