



dCache.ORG

dCache.ORG

dCache

Patrick Fuhrmann



additional funding, support or contributions by





Topics

- dCache distribution channels
- Project news
- SRM 2.2 interoperability
- xRoot protocol interoperability
- dCache 1.8 and SRM 2.2 Space Management
- dCache 1.8 deployment
- Work in progress
- Support issues



Scope of this presentation

This presentation neither contains technical details nor does it advertise features in dCache. Please refer to www.dCache.org for more information.

It is an update on the information provided during the previous LHCC comp. review.

The document concentrates on the deployment of SRM 2.2 in dCache and the upcoming support questions for the near future.



Project news

The Nordic Data Grid Facility (NDGF) is now part of the dCache development team and provides input on issues arising from the fact that the NDGF production dCache spans 4 countries and has to drive different Tape Storage Systems at the same time in different locations.



dCache distribution

dCache is distributed and configured through YAIM for the Tier II's in Europa.

dCache is distributed and configured through VDT for OSG supported Tier II's

dCache is distributed through dCache.org for the Tier I's and configured manually due to the complexity of the Tier I setups.

dCache.ORG

dCache.ORG



SRM 2.2 interoperability

dCache 1.8 has been “Certified” by Flavia's S2 test system to be fully compliant with the WLCG SRM 2.2 specification.

During the PPS test phase some misinterpretations of the agreements on how to handle spaces have been detected and fixed.

The evaluation of the PPS systems by the experiments have been “amazingly” successful. Results have been presented during the pre GDB and the GSSD meetings.

Weekly phone conferences with the dCache Tier I representatives as well as phone conferences with the WLCG, Castor/DPM and dCache management have been extremely helpful. (Many thanks to Jamie and Flavia).



dCache 1.7 as well as 1.8 provides a native xRoot interface.

A support agreement between dCache.org and ALICE is in preparation.



dCache 1.8 and SRM 2.2 Space Management

Details on dCache 1.8 and the SRM protocol versions :

- To use SRM 2.2 in dCache, the upgrade to dCache 1.8 is mandatory.
- dCache 1.8 still supports the SRM 1.1 protocol.
- The use of SRM 2.2 and space management are optional in dCache 1.8
- The decision on *if and when* to enable the SRM 2.2 protocol and space management is up to the sites and the experiments they host.

Details on the dCache 1.8 upgrade procedure :

- The upgrade to dCache 1.8 is a normal rpm software upgrade.
- No data or meta data conversion needs to be done.
- The upgrade can be performed in parallel on all dCache components.



dCache 1.8 (including SRM 2.2) deployment

The Tier I production system deployment started Oct 29 with NDFG and is ongoing. dCache is currently installed at 9 Tier I's.

NDGF	Oct 29	OK
GridKa	Nov 6	OK
SARA	Nov 19	
IN2P3	Nov 26	
FermiLab	Nov 29	
RAL	Dec 3	
Triumf	Dec 3	
BNL	Dec 10	
Pic	Dec 17	

dCache 1.8 will replace dCache 1.7 in the VDT package presumably mid of February

The about 40 dCache Tier II's (LCG & OSG) may join at any time.



dCache 1.8 (including SRM 2.2) deployment (cont.)

Issues detected during the NDGF upgrade :

- › NDGF has not been enabled the space management subsystem yet. (Site decision)
- › This will happen very soon.
- › All issues detected during the upgrade phase have been fixed.

Issues detected during the gridKa upgrade :

- › GridKa has been enabled the Space Management subsystem.
- › There is still one remaining software issue in the space management subsystem, which is currently being worked on. (Problem has been fixed Saturday and will be made available this week.)
- › Due to the fact that based on the WLCG SRM 2.2 agreement, only incoming data is supposed to be managed by the space subsystem, it is not yet clear how to treat data being restored from tape. This is primarily a policy decision. (Should those datasets go to managed space or not?)

Up to now, all upgrades have been done within the scheduled time slot.

There have been no problems detected yet, which prevents us from suggesting to proceed with the scheduled deployment procedure.



Work in progress

Improved monitoring and automated problem detection.

Chimera, the improved dCache file system engine, is currently being tested at various places and will become available beginning of next year. OSG will start investigation mid of December.

We are in the process of implementing the NFS 4.1 interface in dCache.

- › Will very likely become an industry standard.
- › Security is mandatory.
- › No HEP specific data access protocols needed any more (dCap,rfio,xRoot).
- › The OS clients are delivered and maintained by the OS vendors.
- › May be used for local and wide area data access.

We would encourage other SE providers to put efforts in NFS4.1 as well.



dCache.ORG

dCache.ORG

And now for something completely different





Statements by the software providers on general support

FermiLab (quotation)

Support lifetime

Fermilab intends to continue to support SRM 2.2 (on dCache) in the form of bug fixes as long as there is a reasonable HEP community using the software. Work on Future versions of the SRM depends on Fermilab's programmatic needs, and its ability to support the envisioned deployment community.

Provided support for SRM 2.2 (on dCache 1.8)

The Fermilab SRM development team will fix bugs in the SRM v2.2 software and define and clarify behavior that may not be clear in the documentation, and will participate in troubleshooting in cases where it is clear that SRM/dCache software developers are required to resolve an issue.

Support model

It is difficult or impossible to fully test SRM/dCache under real loads and usage patterns. We rely on initial sites to help us commission and identify issues, and in turn for this contribution these sites will receive special consideration.

We view 3 levels of steady-state support:

- Level 1 support: On site expertise, hardware selection and installation, hardware maintenance and support, SRM/dCache installation and configuration, SRM/dCache maintenance and operation troubleshooting.
- Level 2 support: Aggregate knowledge-base of SRM/dCache and associated tools. Provide second level of installation and configuration support and advice, configuration and operation troubleshooting (in particular between sites included), filtering of issues and issuing software bug tickets to dCache/SRM support.
- Level 3 support: Bug fixes, behavior and documentation clarifications, software releases (development)

With the exception of the base Fermilab program and other organizations that fund Fermilab SRM development and support, Fermilab will only provide Level 3 support to the general community. An exception is the commissioning of SRM/dCache 2.2/1.8 at the LHC Tier-1 sites, where developer involvement is expected at Level 2 in return for commissioning the software.

More information on how Fermilab envisions to solve future support is available on request.



Statements by the software providers on general support

NDGF (quotation)

NDGF currently considers dCache the best solution to our storage needs. We will continue to contribute to dCache development while dCache is used at NDGF.

- > We provide first level support for sites deploying dCache in the Nordic countries.
- > NDGF will provide second level support for our contributions to dCache as long as there is a need and we are able to do so, which is at least for as long as we use dCache ourselves.



Statements by the software providers on general support

DESY

Within the dCache community, DESY is responsible for

- › dCache core components
- › The dCache.org infrastructure (technical lead, web, software repository, e-mail lists, ticket system)

We will maintain the dcache.org infrastructure whilst there is a reasonably large community using the dCache technology.

Whilst dCache software is in use at DESY, support will be at least "best effort" for those components contributed by DESY. The availability of additional funding will allow DESY to provide an enhanced level of support. This additional support will be dictated by funding sources, appropriate to level of funding made available and guaranteed for the duration of this funding.(EGEE, d-grid, ...)



Missing support link.

Patrick's opinion

- › To my opinion, the storage and transport of data (data grid) is crucial for the success of the LHC challenge. (Trivial statement)
- › I regard the suggestions made during the GSSD Data Management workshop by Flavia, on how to handle data management support in LCG, as inevitable though not sufficient.
- › It would be careless to leave the management and operation of this area to a group of volunteers.
- › I fear that if there is not a group of experts, dedicated to support the data grid infrastructure, the enormous work, the software providers as well as the site admins have put into the development and installation of those systems has been wasted.
- › A dedicated group of data management experts need to be installed having detailed knowledge of the data management software components, like the Storage Element implementations, the FTS system and the data catalogs as well as their interactions.
- › dCache.org would happily support such a group and would treat any trouble tickets submitted by this group highest priority.



Further reading

www.dCache.ORG

