

eGEE

Enabling Grids for
E-science in Europe

www.eu-egee.org



Using SRM: DPM and dCache

G.Donvito, V.Spinoso
INFN Bari



EGEE is a project funded by the European Union under contract IST-2003-508833

- **Installation**
- **Test set-up**
- **Tests**
- **Using SRM: considerations**
Future tests
- **Conclusions**



DPM Installation



- **Manual Installation:**
 - DPM-Admin-Guide is very useful and almost complete.
 - Expanding the troubleshooting section with more known “trouble”, may help!
- **YAIM Installation:**
 - Almost automatic: small manual configuration is still needed
 - **But...** YAIM does not support, natively, more than one SE per site (a dirty trick is needed)



dCache Installation



- **Manual Installation:**
 - The installation guide is sufficient for standard installation
 - The integration in LCG infrastructure is not so easy
 - Expanding the documentation with some examples for more advanced configuration, may help.
- **YAIM Installation:**
 - The integration in the LCG infrastructure is simpler
 - Some steps need to be done by hand
 - **And...** YAIM does not support natively more than one SE per site (a dirty trick is needed)



Test set up: DPM



- **Server configuration:**
 - **Two machine:**
 - One for: DPM server, DPNS server, MySQL server, gridftp and rfiio
 - One for: gridftp and rfiio
- **Client configuration:**
 - **Two machine with: srmcp client (from dCache), DPM rfiio client**
 - **Other SRM server (other DPM, dCache and CASTOR)**



Test set up : dCache



- **Server configuration:**
 - **Two machine:**
 - One Admin node: Postgres Server, dCache core, dCache opt (dcap, gridftp doors), dCache pool
 - One for: dCache opt (dcap, gridftp doors), dCache pool
- **Client configuration:**
 - Two machine with: srmcp client (from dCache), dccp client
 - Other SRM server (other DPM, dCache and CASTOR)

- **Functionality tests:**
 - SRM: (put, get, copy)
 - Local protocol: (dcap, rfio)
 - Using some deep configurations
- **Stress tests:**
 - Performance
 - Load balancing
- **Fault tolerance:**
 - Daemon crash
 - Hardware crash



Tests with DPM: Functionality tests



- Directory listing on DPM server via SRM: **SUCCESSFUL**
- Copy local fs -> DPM server via SRM: **SUCCESSFUL**
- Directory listing on DPM server via SRM, verifying the presence of the copied file: **SUCCESSFUL**
- Multiple access to the same file by many applications (deleting while transferring file): **SUCCESSFUL**
- Copy local fs -> DPM server via SRM, into a nonexistent directory: **SUCCESSFUL**
- Copy DPM server -> (same) DPM server via SRM: **FAILED**
- dpm-addfs, putting, listing and deleting a file: **SUCCESSFUL**
- Setting RDONLY on a fs (dpm-modifyfs): **PARTLY SUCCESSFUL**



Tests with DPM: Functionality tests (2)



- dpm daemon sync: **FAILED**
- Getting files stored on a removed DPM fs (dpm-rmfs): **FAILED**
- Interrupting transfer of a big file (9GB): **FAILED**
 - Write a file bigger than free space: **FAILED**
- File access through many protocols: **SUCCESSFUL**
- Manipulating fs and permissions: **SUCCESSFUL**
- Remote copy, interaction among many SRM: **PARTLY SUCCESSFUL**



Tests with DPM: Stress tests



- **Balancing multiple requests to different server: SUCCESSFUL**
- **Performance on the single transfer: SUCCESSFUL**
- **Load on the machine during transfer: SUCCESSFUL**



Tests with DPM: Fault Tolerance



- Crashing mysql while PreparetoPut: **SUCCESSFUL**
- Crashing a file server: **SUCCESSFUL**
- Crashing one machine: **SUCCESSFUL**



Tests with dCache: Functionality tests



- Directory listing on dCache server via SRM: **SUCCESSFUL**
- Copy local fs -> dCache server via SRM: **SUCCESSFUL**
- Directory listing on dCache server via SRM, verifying the presence of the copied file: **SUCCESSFUL**
- Multiple access to the same file by many applications (deleting while transferring file): **SUCCESSFUL**
- Copy local fs -> dCache server via SRM, into a nonexistent directory: **SUCCESSFUL**
- Copy dCache server -> same dCache server via SRM: **SUCCESSFUL**
- Getting files stored on a removed pool (killing dcache-pool): **SUCCESSFUL**



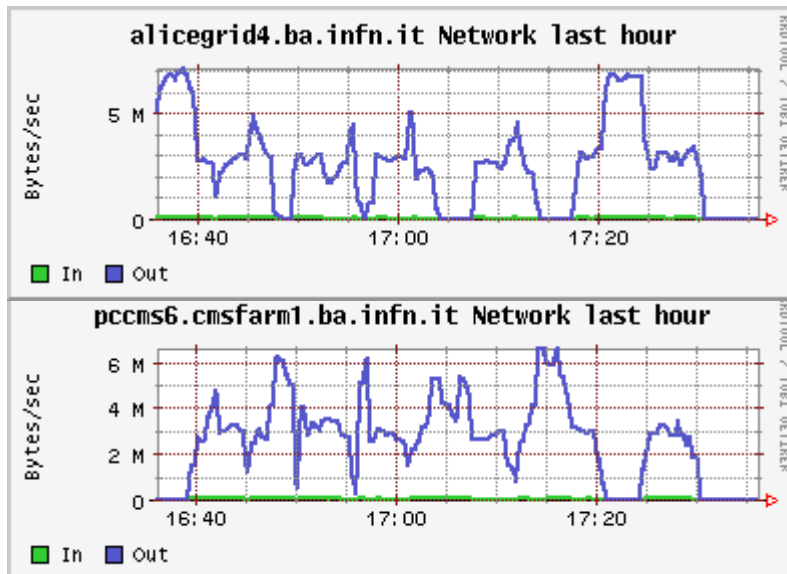
Tests with dCache: Functionality tests (2)



- Interrupting transfer of a big file: **SUCCESSFUL**
- File access through many protocols: **SUCCESSFUL**
- Manipulating fs and permissions: **SUCCESSFUL**
- Remote copy, interaction among many SRM:
SUCCESSFUL
- Multiple requests: **SUCCESSFUL**

Tests with dCache: Stress tests

- Balancing multiple requests to different server: **SUCCESSFUL**
- Performance on the single transfer: **PARTLY SUCCESSFUL**
- Load on the machine during transfer: **PARTLY SUCCESSFUL**





Tests with dCache: Fault Tolerance



- Crashing Postgres while PreparetoPut: **SUCCESSFUL**
- Crashing a file server: **SUCCESSFUL**
- Crashing one machine: **SUCCESSFUL**



Using DPM



- **PRO:**
 - Very simple to install, configure and manage
 - LCG integration quite automatic and up-to-date
 - Very good performance (also with single transfer)
 - Quite little overhead on the servers
 - Support SRM-v2
- **CONS:**
 - Sometimes strange behavior:
 - daemon crashes (dpm-gridftp)
 - DPM not up-to-date to configuration: need to be restarted after some commands
 - srmcp not yet supported
 - Some advanced features not yet implemented (replica of “important” file, load balancing between accepting gridftp connections and writing files...)
 - Scalability not proved
 - Problem when server are not in the same domain



Using dCache



- **PRO:**

- Many advanced features (replication of files, load balancing between accepting gridftp connections and writing files...)
- More robust behavior in production environment
- Support for srmcp (between all known SRM servers)
- Larger scalability (proved in many Tier1/2)

- **CONS:**

- Less “brilliant” performance on single transfer (**but** using more processes the bandwidth can be easily filled)
- More CPU load for copy processes
- Integration in LCG less automatic and up-to-date
- Support only SRM-v1.1



Future tests with DPM



- **Test the new releases as they come out**
- **Test the capability of using DPM with ROOT file access**
- **Test the XROOT access to DPM**
- **Test and use of the DPM installation as back-end for gLite IO-Server**
- **Test deeply the API for accessing files, comparing GFAL and RFIO protocols, in some user's applications**



Future tests with dCache



- **Test the new release**
- **Test the possibility to have more than one replica of (some) files for guarantee the availability and the performance**
- **Test the use of the WN's file system for scratch area for "volatile" files or as space to replicate "frequently accessed files"**
- **Test and use the "production" installation as back-end for gLite IO-Server**
- **Test deeply the API for accessing files, comparing GFAL and DCAP protocols, in some user's applications**

- **DPM is the simplest SRM solution for small Tier 2 sites with no need for advanced features**
 - It is possible to migrate “old” classic SE to SRM with no loss of data in a very easy way
 - **But** ... seems that in a real production environment it needs some attention to fix the problem that can arise
- **dCache seems a good solution for medium-big Tier 2 site**
 - Many advanced features are available
 - The installation is more simple that in the past
 - **But** ... some work need to simplify the integration with LCG