



CLUSTER OF RESEARCH INFRASTRUCTURES
FOR SYNERGIES IN PHYSICS



CERN Development Plan

Ivan Calvet - CERN
CRISP 2nd Annual meeting
18th March 2013





WLCG

- Computing grid mainly composed of 3 grids: EGI, OSG and NDGF
- Collaboration of 170 computing centers
- ~25 Petabytes / year to manage



WLCG
Worldwide LHC Computing Grid



Disk Pool Manager

- Lightweight and grid-aware solution for disk storage management
- Easy to install and manage
- Widely deployed (~200 sites, ~45PBytes)
- Use standard protocols and standard building blocks





DPM: Dmlite

- DPM use a plugin-based library called Dmlite
- Plugins: MySQL, Oracle, Memcache, HDFS, S3, ...
- Supports many frontends: HTTP/DAV, GridFTP, Xrootd, NFS 4.1/pNFS, SRM...





DPM: Virtual File System

- Development plan: Integration of cluster file systems with wide area data storage
- Virtual File System plugin for Dmlite
- Allow:
 - Integration of cluster file systems into existing storage elements from the grid infrastructure
 - Consolidation of the storage of some sites



DPM: 2 VFS

- Two main cluster file systems:
 - Lustre: high performances capabilities, open licensing, 60% of the top 100 supercomputers including Titan (N° 1)
 - GPFS: high performances too, developed by IBM, used for many supercomputers too
- Both very interesting to be used
- Common posix-like commands allows to have a common VFS plugin



Glue 2.0

- Standard Information model use in WLCG
- New client tool for Glue 2.0 called Ginfo (V1.0)
- Query information in a Glue 2.0 BDII by using LDAP requests
- Evolution of the client depending from feedback



Summary

- DPM is a lightweight and grid-aware solution for data storage
- Uses a plugin-based library
- VFS plugin will be developed for cluster file systems
- Glue 2.0 has a new client tool