

Data Management

Stephen Burke, PPARC/RAL

Jeff Templon, NIKHEF





Overview



Current situation + some outstanding problems:

- ◆ GDMP status - Jeff Templon
- ◆ Directing jobs to data
- ◆ Replica Manager
- ◆ Mass storage



GDMP Status

Jeff Templon



Sending jobs to data



- ◆ The JDL can specify a list of LFNs and an access protocol, and the address of the Replica Catalogue.
- ◆ Protocols are "file" for local (NFS) access, "gridftp" and RFIO - but Castor files can't easily appear in an RC so RFIO is not so useful (see later).
- ◆ Currently all SEs support all protocols.
- ◆ The job will go to a CE which has all files locally.
- ◆ Information about the local system and filenames is available through the BrokerInfo interface.



Replica Manager



- ◆ No real Replica Manager yet (i.e. no code to replicate files automatically), but some prototype commands are available.
- ◆ There are commands to copy a file into or between SEs, and to update the Replica Catalogue.
- ◆ There are commands to delete files and RC entries .
- ◆ Designed for single files rather than batches.
- ◆ Does not write full file information into the RC.
- ◆ See the tutorial by Heinz Stockinger.



Mass Storage: RFIO



- ◆ The basic solution for accessing tape stores is RFIO. Files can be copied between Castor and disks at CERN using *rfcp* as normal.
- ◆ Anyone with a CERN tape account is mapped to that account, otherwise to generic accounts with enough privilege to read and write.
- ◆ *rfcp* can be used from a WN at CERN. It is not currently on the UIs.
- ◆ This usage can be extended to HPSS at Lyon.
- ◆ RFIO can do direct file access with API calls.



RFIO drawbacks



- ◆ RFIO only works on machines local to the tape store (security issues).
- ◆ Files in Castor can't easily be recorded in a Replica Catalogue, and even if they could the WP2 tools would not know how to deal with them.
- ◆ There is no GridFTP access to tape files (coming soon?).
- ◆ RAL has a completely different system (virtual tapes), although this can also be used in a similar way.



GDMP hooks



- ◆ GDMP has some hooks to allow MSS access. When a file is replicated from another site to CERN it can be automatically copied to a fixed location in Castor.
- However, locally-created files are not copied automatically.
- ◆ When GDMP gets a request to replicate a CERN file, if the file is not on disk it tries to stage it first (this currently fails due to a bug).
- Again, local access does nothing.



WP5 staging scripts



- ◆ The current testbed has some staging scripts which use rfcop to copy files to/from a fixed directory in Castor.
- ◆ Drawbacks: file names can clash, Castor has a maximum file name length of 130 (?) characters, the scripts are not callable by users.
- ◆ New scripts will solve all these problems, at the cost of using a relatively opaque file naming scheme (MD5 hashes).



MSS Problems



- ◆ No disk space management, users need to manage the space themselves.
- ◆ Automatic staging is limited to replication, in general users have to stage in and out themselves.
- ◆ Files are staged to a specific place with opaque filenames, which may be hard to integrate with an external file catalogue.
- ◆ Files in other places in Castor cannot be registered in Replica Catalogues or manipulated by Grid tools.
- ◆ Experiments need to define their needs to see if they can live with this in TB1.2.