

ENABLING ACCESS TO MASS STORAGE SYSTEMS

Dr John Gordon, Dr Jens G Jensen, Owen Syngé
GridPP and e-Science departments, Rutherford Appleton Laboratory

Key words to describe the work: EU DataGrid, Grid storage, Grid protocols, mass storage, data transfer, data security

Key Objectives: Our aim is to build and deploy a flexible and robust interface to a large range of mass storage systems. This work is being done in the framework of the EU DataGrid so it will solve Grid storage requirements. The flexibility of the system allows it to be deployed also outside a Grid framework.

Motivation for the work (problems addressed): A number of different storage systems are in use in the scientific communities. We present a solution for accessing data in a uniform manner and for replicating it across sites. This work will be of interest to anyone with large scale storage requirements, e.g. the UK Grid community.

Storage requirements in the scientific community are serviced by disparate storage systems, ranging from the simple disk system to large systems with tapes and robots. As scientific computation moves towards a Grid environment, a need arises for accessing those storage systems through Grid protocols. Some files in the storage systems will then be made available to users on the Grid, others will not. Conversely, Grid users will be able to write files into the storage systems, and the files may later be replicated to other storage systems. Of course, access to the storage will be granted to authorised users only, and access is auditable by storage administrators.

To address these problems, the European Union DataGrid project is building the "Storage Element", a robust and extensible interface to mass storage systems. This work is being done by the GridPP and e-Science groups at RAL.

We have designed an SE software system. The role of the SE is primarily to sit between the client and the Mass Storage System (MSS); to hide the MSS differences from the client and to allow access to the MSS using protocols that it does not naturally support. In addition to this role, the SE will also provide other Grid functions as applied to data access. For example: security; access control; quotas; monitoring; logging; network transfer estimation.

To the outside world, the SE will provide three functions:

- For data transfer, it will support existing protocols such as RFIO, and GridFTP, and will be extensible to new protocols that may appear. It will allow these protocols to access the MSS.

There will also be an API to access files in an SE and replicate files between SEs.

- For control, it will provide a range of functions such as reservation, pinning, deletion, and transfer time estimation. An API to these functions will be defined and implemented.
- For information, it will act as information providers to the DataGrid Information Service, providing metadata about the SE, the underlying MSS, and the files therein.

The design of the SE follows a layered model with a central core handling all paths between client and MSS. This approach was chosen as simpler to implement, for multiple client/MSS combinations, than implementing each combination separately, as the latter would lead to combinatorial inflation with resultant maintenance problems.

The SE that will be implemented from this design can be extended to support new protocols for data transfer, new mass storage systems, new Information Services, and new Grid Services that wish access to data.

The SE will be of interest not only to scientific communities but for anyone who wishes to share large file data access within an institute or across international collaborations. The SE enables anyone with a Public Key Infrastructure (PKI) to securely share files between existing mass storage systems using common Grid tools. In addition, the architecture of the SE is flexible enough to enable people with a storage system that is not supported to add support for this system into the SE. This flexibility also means that, although the SE is being built for a Grid environment, people who do not

have a PKI or do not need Grid security can easily modify the interfaces to accept HTTP or FTP and then use the SE to access their mass storage systems through these protocols.

The SE has a liberal open source licence which allows anyone access to the source code and still allows the SE to be modified or used in proprietary products.