# Experiences with supporting mass storage for multiple communities

Submitted as Paper and mini-presentation to:
Data Integration in eScience: what does the Grid give us?

**Authors:** P Berrisford, D Corney, T Folkes, J Jensen, J Mencak, D Ross, B Strong, O Synge, CCLRC

## Abstract

Access to reliable high capacity and high performance storage is becoming increasingly important to many scientific projects.

Rutherford Appleton Laboratory runs a Petabyte-scale mass storage facility [ADS] which serves a wide range of scientific communities, from within the research council, from other research councils, and other communities. This paper describes our experiences with providing storage services for all these scientific communities.

Some communities use Grids to analyse their data, others use traditional clustered computing, and others want to backup or archive their data. Some use Java on Microsoft Windows or portals, others use C on Unix, or scripting languages. Some communities have clients that talk directly to mass storage, others require higher level services in addition to the tape service. They have different access patterns – backup data is usually not read back, some analysis data is read as soon as it's written. Given this huge range of requirements, we ask the question whether it is feasible for a single facility to offer this width of services, and, if not, how to decide who we can support and who we cannot.

This paper will be of interest to anyone working with storage and data management issues, from the Grid or otherwise. Another important aim of this paper is to communicate the issues and problems to the scientific communities, in order to encourage a dialogue which will help us improve the services. The paper will be readable by non-experts.

## Discussion

Different communities have different tools and software infrastructures. The main focus of this paper is to discuss our experiences with supporting multiple communities and interfaces to the ADS mass storage system, briefly covering the advantages and disadvantages of each interface, and discussing whether, and how, they can be integrated. In particular, we will cover Grid interfaces such as the Storage Resource Manager [SRM] and Storage Resource Broker [SRB]. We shall see how "thick" interfaces can enhance the usability of the ADS (e.g. by adding Virtual Organisation management), and, conversely, the ADS improves the service provided by the interface. We then look into the architecture of the underlying system that supports the interfaces; in particular, how the interfaces help isolate users from architecture changes and when they don't, and what we can do when they don't. We also look at the overhead imposed by the interfaces, e.g. the time for the client to start up and negotiate tape space.

As specific examples, we describe lessons learned in providing services for BBSRC [BBSRC] and LHC [LHC]. There are many contrasting differences between the requirements of the two groups, such as access patterns (write for backup or write for immediate analysis), technology (implementing current technology or the need to move forward), security, writing small files or very large files (mass storage systems don't handle small files well), and the required guarantee of reliability.

We discuss the advantages and disadvantages of supporting Web Services interfaces.

We also look at the need for sharing files between different communities or different VOs within a community, and accessing files via different interfaces: to which extent this is possible, and how it has been, or can be, achieved. This problem can occur not just when there is a need to share files, but also when migrating from one version of an interface to another, or in a disaster recovery. A central aspects of this problem is how to maintain the link between the data and the metadata, or between one version of the metadata and another, and to ensure that they are consistent. We describe how the standardisations of metadata and schemata facilitate data integration.

One of the currently open issues is that the requirements of the Particle Physics community [GridPP,Tier1] such as LHC drive the process forward – not just storage capacity, but also networks, interfaces and tools and other software, etc. The other, non-PP, communities often benefit from this process. The converse doesn't seem to happen. We look at whether there are cases where the PP community can benefit from the experiences of the non-PP communities, and ask how these cases can be communicated effectively back to PP, and what we as storage providers can do to encourage this dialogue. We hope this paper will help stimulate discussion.

**References**


ADS – http://www.e-science.clrc.ac.uk/web/projects/petabyte

BBSRC – http://www.bbsrc.ac.uk/

GridPP – http://www.gridpp.ac.uk/

LHC – http://lcg.web.cern.ch/LCG/

SRB – http://www.npaci.edu/DICE/SRB/

SRM – http://sdm.lbl.gov/srm-wg/

Tier1 – http://www.gridpp.ac.uk/tier1a/