

Storage Resource Management for Data Grid Applications

PI: Arie Shoshani (LBNL)

CoPIs: Don Petravick (Fermilab), Alex Sim (LBNL)

Other staff: Junmin Gu , Vijaya Natarajan, Alex Romosan (LBNL),

Timur Perlmotov (Fermilab)

(<http://sdm.lbl.gov/srm>)

Summary

Storage Resource Managers (SRMs) are middleware components whose function is to provide dynamic space allocation and file management on shared storage components on the Grid. They complement Compute Resource Managers in providing storage reservation and dynamic information on storage availability for data movement, and for planning and execution of Grid jobs. Over the last two years, we have developed the concepts and functionality specification of SRMs, and implemented SRMs at LBNL and Fermilab. This technology is already being used by SciDAC supported projects. This includes the STAR, CLAS, CDF, and Lattice QCD experiments as part of the Particle Physics Data Grid (PPDG) project, and the Earth Science Grid (ESG) project. Furthermore, we have coordinated an international effort to standardize SRM interfaces as Web services. A GGF Working Group was established for this purpose.

Managing shared storage resources on the Grid is a necessary and complex task because of the diversity of the storage resources. Storage resources can vary in complexity: from a single disk to large sets of disk caches or disk RAIDs, or mass storage systems (such as HPSS) that provide access to data on robotic tape systems. Making such resources sharable through Grid middleware technology requires that these systems are exposed through a uniform interface. Thus, requesting space from any of these systems should look the same to a client. We have shown in this project that through international collaboration of achieving agreements on such common interfaces, various storage systems can interoperate. Most notably, SRMs have been built for several mass storage systems both in the US and in Europe and shown to interoperate smoothly. Furthermore, several SRMs were built for disk systems as well, and they interoperate with SRMs for mass storage systems. This approach of standardizing on the functionality and the

interfaces of SRMs is the backbone to our success. It allows multiple groups to implement their own SRMs and thus make the underlying storage system viewed as a Grid service.

Running jobs on the grid requires the coordination of multiple middleware components. Typically, an application makes a request for running a job to the “request manager”. The request manager may include a component called a “request planner” that figures out the best way to run the job by consulting metadata catalogs, file replica catalogs, monitoring information, etc. The plan can then be handed to a “request executor” that will then contact compute resource managers and storage resource managers to allocate resources for executing the job. The SRMs are components that a “request executor” can rely on to perform space allocation and file management when requested to do so.

We highlight below some of the main achievements of this project so far.

- Initiated the concept design and functionality of SRMs.
- Coordinated the specification of a common SRM interface, SRM version 1, with the participation of Fermilab and Thomas Jefferson National Accelerator Laboratory (TJNAF). This led to the Implementation of SRMs at Fermilab on top of the Enstore Tape Storage System coupled with the dCache distributed disk caching system., at TJNAF on top of JASMine mass storage system, and at LBNL on top of HPSS mass storage system.
- Coordinated the acceptance of SRM version 1 by the European Data Grid project. This led to development of an SRM on top of the Castor mass storage System at CERN.
- Developed an SRM for a UNIX-based disk system, called Disk Resource Manager (DRM). This is installed and used at LLNL and NCAR as part of the ESG project.
- The SRM for HPSS developed at LBNL, called Hierarchical Resource Manager (HRM-HPSS), was adapted and deployed in various environments at LBNL/NERSC, ORNL, and BNL for use by the PPDG and ESG projects.
- HRM-HPSS was modified to work with a legacy mass storage system at NCAR, called HRM-MSS. This was deployed at NCAR for the ESG project.
- Deployed SRM developed at Fermilab to the High Energy Physics Lattice QCD and CMS experiments.
- Developed a tool for robust massive data replication for entire directories of thousands of files, called DataMover. See Figure 1. It is deployed by the Nuclear Physics STAR experiment for routine production replication, as well as ESG.
- Coordinated the specification of version 2 of SRMs that is the basis of the next version of SRMs to be developed.

Institutions involved were Fermilab, TJNAF, CERN and Rutherford Lab in England. This version will support negotiated space reservations and directory service.

- Initiated a GGF Working Group with the goal of standardizing SRM web services.

Plans for future work include four major thrusts:

- The development of SRM version 2 by LBNL and Fermilab. This requires a major enhancements in the capabilities of SRMs to support negotiated space reservation, and directory management.
- The development of GGF-based SRM standards for the Grid community. This involves the coordination of multiple institutions, and getting industry involvement.
- The deployment of SRM software in additional SciDAC projects in various scientific communities.
- Enhance SRM design to support the interaction with other middleware components, such replica catalogs, virtual organization managers, and request planners and executers.

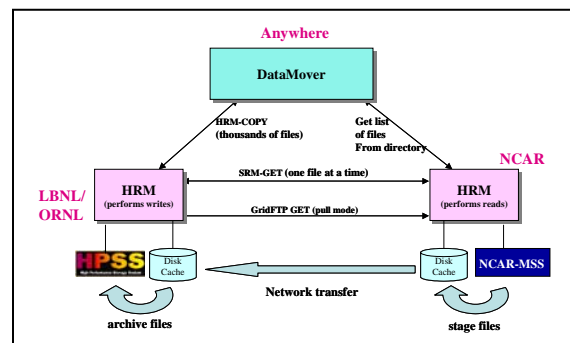


Figure 1: Robust File Replication of entire directories using SRMs

For further information contact:

Dr. Arie Shoshani
 Lawrence Berkeley National Laboratory
 Tel: (510) 486-5171
 Email: shoshani@lbl.gov