Storage Resource Managers (SRM) in the Earth System Grid

Arie Shoshani, Alex Sim, Junmin Gu, Vijaya Natarajan,

Lawrence Berkeley National Laboratory

Summary

The Storage Resource Management technology was developed in response to growing needs of managing large datasets on a variety of storage systems. This technology increases the scientist productivity by eliminating the tedious and time consuming tasks of managing storage, performing robust file movement, and dealing with security requirements at various storage sites. Storage Resource Managers (SRMs) are used for space management, streaming data to the scientists' sites, and removing automatically unused replicated files from shared disk cache storage. The SRM approach is to develop a standard interface that allows multiple implementations by various institutions to interoperate. The SRM used by the Earth Systen Grid (ESG) project, called the Berkeley Storage Manager (BeStMan) was developed by ESG members at LBNL. It is used to interface tansparently to the HPSS mass storage systems at LBNL and ORNL, the MSS mass storage system at NCAR, to disk pools at LANL and LLNL, and to manage the disk cache that serves users at the ESG Gateway in NCAR.

In ESG data files are stored at various sites, on shared disk storage and on mass storage systems (MSSs), such as HPSS. A typical use case is for a user to login to the ESG Gateway (portal), browse the metadata, and identify files of interest to download. The Gateway allocates storage space using in a disk cache managed by BeStMan, and requests the local BeStMan to get the files from local or remote disk cahces and MSSs. The local BeStMan then connects all remote BeStMan systems to get the files. This is shown in Figure 1, where BeStMan Storage Resource Managers (shown in purple) communicate with each other transparently, and use GridFTP to move the files.

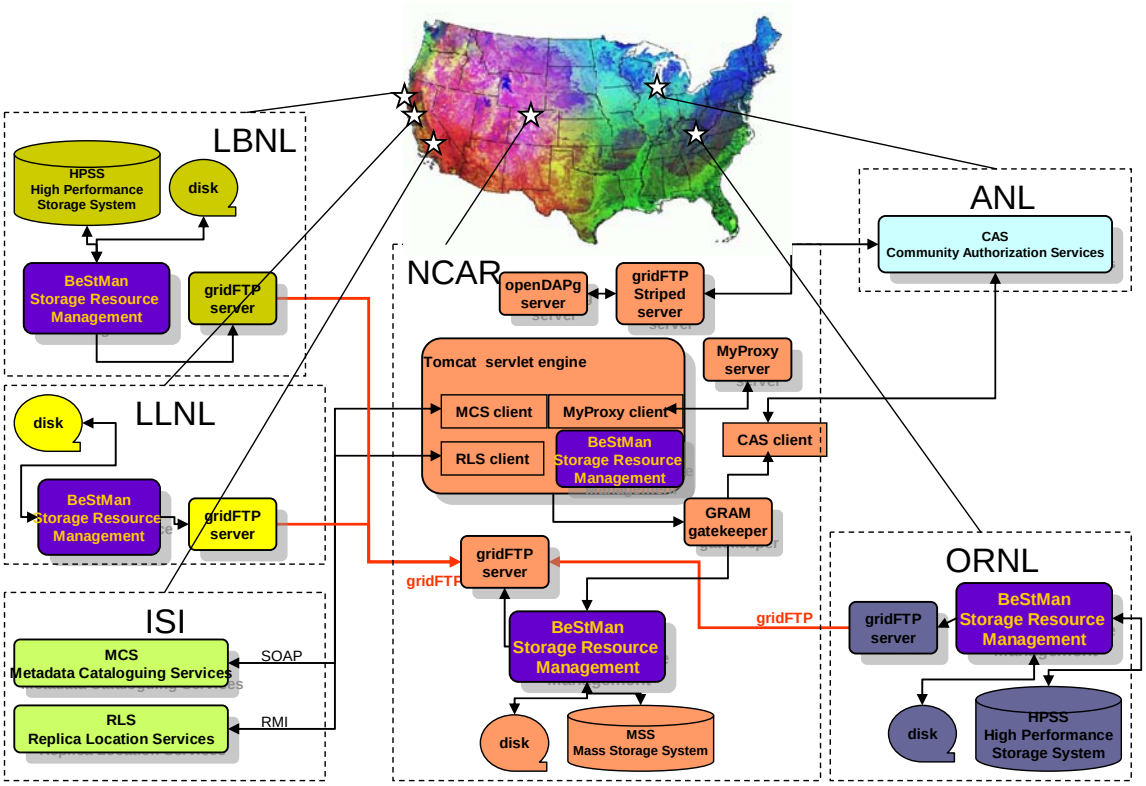

Figure 1: Use of the Berkeley Storage Manager (BeStMan) in ESG

SRM Functionality includes: (i) accepting multi-file request, managing a queue, allocating space; (ii) invoking multiple concurrent file transfers to get/put files from/to remote sites; (iii) pining files for a limited lifetime, accepting early release of files; (iv) queuing staging and archiving to MSSs; (v) providing automatic garbage collection; (vi) monitoring and recovering from failures of file staging, archiving, and transfer.

Once all the requested files are moved locally, they can be served to the users. The users can download the files one by one if only a few files are requested or using specialized tools if a large number of files are request. One of the tools developed for this purpose by the ESG team at LBNL is called DataMover-Lite (DML). DML was a concept that came about because of the necessity to manage multi-file transfers to a client machine. DML is considered "light-weight" in that it can be easily downloaded and used by clients. It can be used to transfer entire directories as a directory-to-directory action.

Figure 2 shows the interface that users can use to monitor progress. At the bottom of the screen information about transfer rate of each file is shown, as well as its size and its source and target URLs. DML is currently in use, and is available from the ESG Gateway. It was implemented for various platforms including a PC and a MAC version. So, far the DML has been used to move files through the data portal only. Currently DML communicates to the gateway to download files using LAHFS based http (Lightweight Authorized HTTP File Server). For future version, it is desirable that DML communicates directly with the SRM at the gateway or directly with SRM or GridFTP servers on data nodes.

For further information on this subject contact:

Name: Arie Shoshani
Organization: Lawrence Berkeley National Laboratory
Email: shoshani@lbl.gov
Phone: (510)485-5171

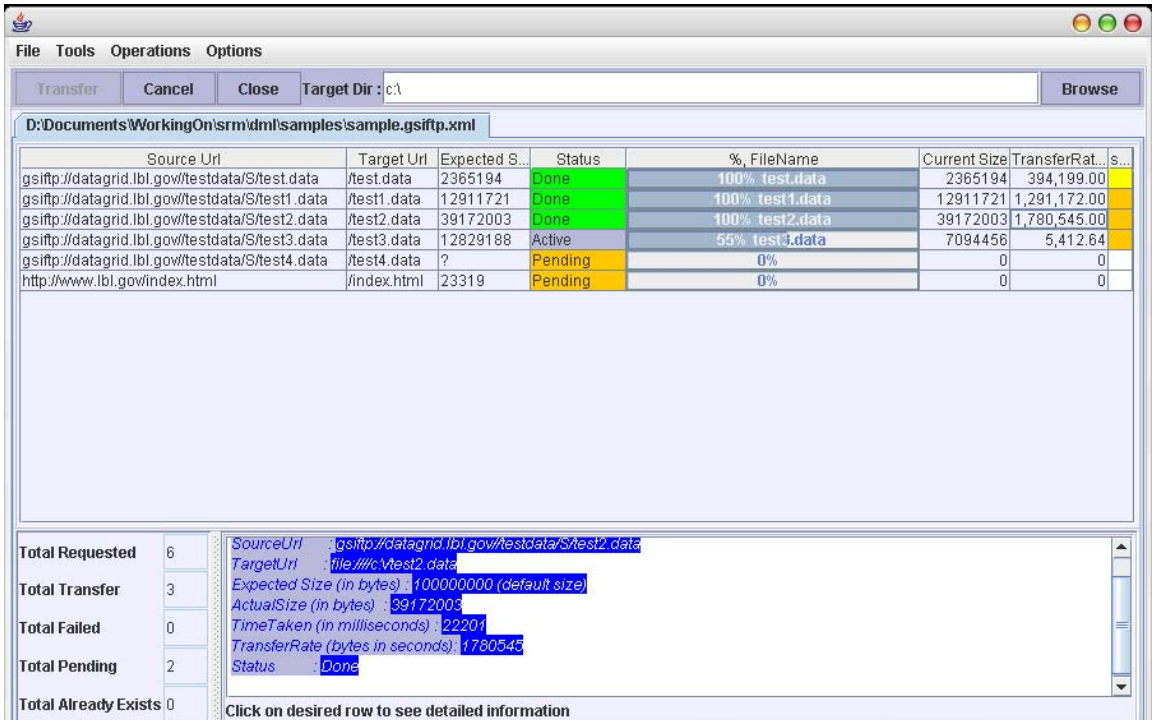

Figure 2: DataMover-Lite interface screen showing progress of multi-file transfer request