

Grid Storage Management Working Group GSM-WG Global Grid Forum, Data Area

Administrative Information	1
Charter	1
<i>Focus/Purpose</i>	<i>1</i>
<i>Scope</i>	<i>Error! Bookmark not defined.</i>
<i>Goals</i>	<i>2</i>
<i>Deliverables/Milestones</i>	<i>3</i>
Management Issues and Evidence of Commitment	3
Exit Strategy	4
References, Pre-existing Documents	5

Administrative Information

Name and Acronym: Grid Storage Management Working Group, GSM-WG.

Chairs:

- Arie Shoshani (Lawrence Berkeley National Laboratory, USA) arie@lbl.gov
- Peter Kunszt (CERN, Switzerland) Peter.Kunszt@cern.ch

Secretary/Webmaster:

- Alex Sim (Lawrence Berkeley National Laboratory, USA) asim@lbl.gov

Email list:

- gsm-wg@gridforum.org

Web page:

- <http://sdm.lbl.gov/gsm>

Charter

Focus/Purpose

We define Storage Resource Managers (SRMs) as middleware components whose function is to provide dynamic space allocation and file management of shared storage components on the Grid. They complement Compute Resource Managers and Network Resource Managers in providing storage reservation and dynamic information on storage availability for the planning and execution of a Grid job. SRMs manage two types of resources: space and files. When managing space, SRMs negotiate space allocation with the requesting client, and/or assign default space quotas. When managing files, SRMs allocate space for files, invoke file transfer services to move

files into the space, pin files for a certain lifetime, release files upon the client's request, and use file replacement policies to optimize the use of the shared space. SRMs can be designed to provide effective sharing of files, by monitoring the activity of shared files, and make dynamic decisions on which files to replace when space is needed. In addition, SRMs perform automatic garbage collection of unused files by removing selected files whose lifetime has expired when space is needed.

In this Working Group we will focus on the definition of the functionality of a standard *SRM interface*. We see this as a short-term effort to have a workable interface that Grid projects could immediately make use of to resolve interoperability issues between storage systems.

Scope

There was a significant coordinated activity so far in the US and Europe on defining the functionality and the interfaces of an initial version of SRMs, mainly for High Energy Physics laboratories.

This resulted in a specification of the *SRM interface* version 1 [R6] that was adopted by several institutions, including several laboratories in the US and in Europe. More significantly, several SRM prototypes to existing Mass Storage Systems (HPSS, Enstore, JASMine, Castor, MSS-NCAR) were developed that adhere to the same *SRM version 1 interface*, and interoperability between them on the Grid was demonstrated [R3]. Some SRMs are already being used in production environments. It is worth noting that this interface is defined using a WSDL interface [R8], thus defining the interface in an OGSA-compliant manner is a straightforward evolution that the GSM-WG will probably decide to pursue.

The same initial group of people has focused its more recent activity on the definition of the specification for *SRM version 2 interface* [R7] which includes dynamic storage resource reservation management and dynamic support of user defined directories over the storage resources that the SRMs manage.

Although the initial SRM requirements came from the High Energy Physics community, the *SRM interface* was designed to be application neutral, and was applied so far to climate modeling and medical applications as well. The main concepts are dynamic management of storage space and files in the storage space, concepts that do not refer to any specific application area. One of the goals of the working group is to expose the design to any application domain including scientific and commercial interests, and verify that general purpose Grid Storage Managements needs are addressed.

We solicit active participation in this effort from all interested parties, both from academia/research and industry. In our experience so far the *SRM interface* was one of the most successful demonstrators of the interoperability efforts between different Grid projects in the US and in Europe.

Goals

We will produce *GGF recommendation documents* for progressive *SRM interface specifications*, based initially upon the work that has been done to date; as well as the description of a certification test suite that will certify whether an implementation truly complies with the given interface recommendation. The documents will contain the detailed specification of the interface, including API, semantic behavior and failure modes. The document will also include a WSDL specification of the interface. The specification document will be complete enough so that any person/organization that is interested in providing an implementation of SRM is able to do so without any extra information.

We foresee two sets of recommendations for a *basic and an advanced SRM interface*. The exact split between these two remains to be discussed but the rough outline is:

Basic	Advanced
Local multi-file access	Remote multi-file access
Storage reservation	Authorization
Directory management	Access control
OGSA compliance	Monitoring

A draft recommendation of the basic *SRM interface* should be available very soon so that those who wish to implement it may still influence its content through iterations in the time to follow. This is realistic based on the existing documents [R7].

Deliverables/Milestones

Basic SRM interface GGF recommendation – outline	03.2004
Advanced SRM interface GGF recommendation – outline	06.2004
Basic SRM interface GGF recommendation – draft	10.2004
Basic SRM interface test suite description – draft	02.2005
Basic SRM interface GGF recommendation – final	06.2005
Basic SRM interface test suite description – final	10.2005
Advanced SRM interface GGF recommendation – draft	02.2006
Advanced SRM interface test suite description – draft	06.2006
Advanced SRM interface GGF recommendation – final	10.2006
Advanced SRM interface test suite description – final	02.2006

Management Issues and Evidence of Commitment

Arie Shoshani has already put considerable effort during the last two years into coordinating the SRM community effort between the HEP laboratories. The same amount of time will be available to him for GSM-WG. Peter Kunszt has a mandate in the context of his current work to work on standardization and interoperability issues in the area of Grid Storage and Data Management. There have been already several face-to-face meetings (at least once per year) during the specification phase of the first and second version of the interface as described by [R6] [R7]. We plan to hold at least one such 2-day workshop per year, open to all committed participants. Based on experience, most of the work will be done at the participating institutes where early prototypes of the SRM specification will be developed and tested.

At the GGF meetings where we plan to have a WG session at least twice a year, we plan to discuss the status and to coordinate the work of the implementing groups. Currently we have the following institutes who provide SRM implementations and are actively involved in the SRM definition process (and we would like to solicit more participants through GGF, especially industrial ones): Lawrence Berkeley National Laboratory (Berkeley CA, USA), Fermi National Accelerator Laboratory (Batavia IL, USA), European Organisation for Nuclear Physics CERN (Geneva, Switzerland), Deutsches Elektronen Synchrotron DESY (Hamburg, Germany), T. Jefferson National Accelerator Facility (Newport VA, USA), CCLRC Rutherford Appleton Laboratory (Oxon, UK). Additional sites that are deploying SRMs are: Brookhaven National

DRAFT
Category: Charter
Data Area

Editors: Arie Shoshani, LBNL
Alex Sim, LBNL
Peter Kunszt, CERN

Laboratory (Brookhaven NY, USA), National Center for Atmospheric Research (Boulder CO, USA), Oak Ridge National Laboratory (Oak Ridge TN, USA), Lawrence Livermore National Laboratory (Livermore CA, USA). All of these laboratories have already invested effort into SRM and potentially plan to do so in the future.

For GSM-WG we will name a liaison officer whose task is to monitor communities (esp. GGF groups) that are relevant to GSM-WG. Input from and to these groups will be channeled and organized through the liaison person. We have identified areas of common interest with the following GGF groups: GRAAP, DAIS, GFS, OREP (see below).

Exit Strategy

The GSM-WG is considered 'done' after having delivered its last deliverable.

Related topics at GGF

The following existing working groups work on related topics and coordination with them is in our interest:

GRAAP Grid Resource Allocation Agreement Protocol Working Group [R9]. Defines advance reservation and other terminologies relevant to GSM-WG. The space reservation functionality of the SRM may be viewed as a special subset of GRAAP: the data store resource scheduling problem. GRAAP has addressed the network and CPU but they state that other resources should be able to follow the specification. By coordinating carefully with GRAAP we could achieve a win-win situation by specifying data resource scheduling through the SRM space reservation mechanism. Issues to be discussed are whether the GRAAP concepts are mature enough to be included in the basic specification or whether these need to be tackled in the advanced specification.

DAIS Database Access and Integration Services Working Group [R10]. Although the DAIS WG works on specification for a Grid Database service, many of the operations through a DB interface address similar issues as SRM operations on files. They have also started to look at file access, which seems to be outside of their original scope, and more inside ours. Since for them it is a new domain to look into, again we could establish a win-win situation by synchronizing the specification for file-based access with DAIS, leveraging our different experiences.

OREP OGSA Data Replication Services Working Group [R11]. Looks at replication services in the OGSA context, especially the Replica Location Service and Replica Management. The services being specified here might need the SRM services for space reservation in order to offer higher-level functionality. Or for the advanced GSM-WG specification the SRM might make use of the OREP services to achieve a more complex replica management functionality.

GFS Grid File System Working Group [R12]. Will provide specifications of the virtual namespace that will allow association of access control mechanisms and meta-data of the underlying physical data sources. The GSM-WG spec is not concerned with *virtualized* hierarchical namespaces, but rather with the *storage system* hierarchical namespaces. There is a need to have a mapping between the *virtual* and the *storage* name spaces. The semantics of the GFS and GSM specifications will have to be carefully coordinated to make this mapping possible. This is also true for the access control semantics defined in GFS (global access control) and SRMSRM (local access control).

With all of these groups GSM-WG will actively pursue a dialog through the designated liaison officer who might enlist help from other members of the GSM-WG.

We see these groups as complementary according to the following table:

Virtual Namespace, Metadata	GFS
File Replication	OREP

Storage and file Management	GSM
Database and file Access	DAIS
Transfer	GridFTP

References, Pre-existing Documents

- [R1] *Arie Shoshani*, Storage Resource Managers: Middleware Components for Grid Storage, MSS 2002 <http://sdm.lbl.gov/srm/documents/srm.mss02.pdf>
- [R2] **Storage Resource Managers: Essential Components for the Grid**, *Arie Shoshani, Alexander Sim, and Junmin Gu*, in Grid Resource Management: State of the Art and Future Trends, Edited by Jarek Nabrzyski, Jennifer M. Schopf, Jan weglarz, Kluwer Academic Publishers, 2003, <http://sdm.lbl.gov/~arie/papers/SRM.book.chapter.pdf>
- [R3] Uniform Grid Access to Different Mass Storage Systems, Joint SRM demo, SC 2002 Nov. 2002 FNAL, JLAB, LBNL <http://sdm.lbl.gov/srm/documents/sc2002/sc02-mutli.mss.ppt>
- [R4] *Alex Sim*, Grid File Replication using Storage Resource Management, GGF 6 Presentation Sep. 2002 <http://sdm.lbl.gov/srm/documents/asim.ggf6.ppt>
- [R5] *Arie Shoshani*, Storage Resource Management, GGF 4 Presentation Feb. 2002 <http://sdm.lbl.gov/srm/documents/02.02.srm.joint.design/index.htm>
- [R6] SRM version 1 specification, <http://sdm.lbl.gov/srm/documents/srm.v1.0.doc>
- [R7] SRM version 2.1 specification, <http://sdm.lbl.gov/srm-wg/doc/SRM.spec.v2.1.final.pdf>
- [R8] WSDL specification of SRM v2.1 <http://home.fnal.gov/~timur/srm/srm.wsdl>
- [R9] Grid Resource Allocation Agreement Protocol Working Group <http://www.fz-juelich.de/zam/RD/coop/ggf/graap/graap-wg.html>
- [R10] Database Access and Integration Services Working Group <http://www.cs.man.ac.uk/grid-db/>
- [R11] OGSA Data Replication Services Working Group <http://www.isi.edu/~annc/orep.html>
- [R12] GFS Grid File System Working Group <https://forge.gridforum.org/projects/gfs-wg/>