

# Storage Interfaces: SRM and Clouds

## 1 Introduction

Storage systems used by HEP require remote management. The availability of a common API reduces the complexity of client-side software. Currently, the SRM v2.2 API is widely used, which provides many advanced features.

There are several problems with relying on SRM v2.2 for management:

- Not all HEP storage is manageable through SRM.
- Not all Storage Element implementations provide the complete implementation of SRM specification.
- Not all of SRM v2.2 specification has proved useful.

Within the commercial sphere, different approaches to managing storage have appeared. Perhaps two of the most successful are Amazon's Simple Storage Service (S3) and DropBox. Both services use proprietary communication for managing storage, with S3 inspiring open-source projects that emulate the interface.

In addition to such proprietary interfaces, standards exist that provide comparable levels of functionality. CDMI is an emerging standard interface for managing storage. In addition, WebDAV provides many of the functionality of SRM, with Microsoft's SkyDrive providing it as the interface to their storage service.

This document breaks down the SRM protocol into basic functional blocks. These descriptions of functionality provide a reference when comparing the functionality of different protocols. It concludes with recommendations with regard to adoption of storage protocols with WLCG community.

## 2 Functional description of SRM

The SRM protocol is specified as a collection of 39 methods. Although not specified, all implementations provide access to these methods using a common mapping to SOAP with HTTP transport.

Rather than describing each of the SRM methods, this section describes the functionality that the SRM protocol provides in terms of the effect on the storage system; an item may encompass multiple SRM operations. This description is grouped based on five high level concepts.

### 2.1 Transfer Management

This section includes behaviour that affects how data is transferred. The following areas have been identified:

- Upload / download a complete file
- Queuing transfers.
- Suspend/resume transfers.
- The ability to cancel a queued transfer.
- Resource provisioning (uploading useful data)
- Permit SE to load-balance client requests over multiple transfer end-points
- Back-pressure (SE tells client to slow down)

- Trigger third-party copy
- Cancel requested third-party copy
- Negotiating a transport protocol that client and server support.

## **2.2 Namespace interaction**

This section includes functionality that adjusts the namespace of a storage element

- Querying information about a file (stat)
- Allowing uploading of data integrity information
- Deleting data and directories
- Changing ownership
- Changing permissions and ACLs

## **2.3 Storage Capacity Management**

Concept of guaranteeing availability of storage capacity.

- Ability to query how much capacity is being used (like df)
- Create new reservations
- Assigning characteristics to a reservation, like access latency and retention policy.
- Targeting uploads to specific reservation
- Moving files between reservations
- Removing reservations---empty or containing files.

## **2.4 File Locality Management**

Controlling the movement of data from tape to disk.

- Non-blocking trigger that files should be staged from tape
- Guarantee that a file will not be garbage collected for a finite period
- Allow ad hoc decisions when a file is eligible for garbage collected (cancel pin)

## **2.5 Server Identification**

These is functionality about the server that is independent of its ability to store data.

- Allow simple test of service availability
- Allow discovery of arbitrary key/value information

# **3 Comparison with alternative solutions**

This section provides a comparison between SRM protocol and that of alternative systems.

## **3.1 Amazon S3**

From commercial service; open-source implementations that emulate interface.

### **3.2 CDMI**

Strong interest from EGI. DDN advertise[1] that they fully support CDMI. NetApp are also very active in this area. OpenStack expected to add support for CDMI (as “technological preview”) in a few months.

[1]

<http://www.storagebytesnow.com/2011/10/13/datadirect-networks-is-first-with-full-cdmi-support-2/>

### **3.3 WebDAV**

Used by Microsoft's SkyDrive service. Clients and servers commonly available.

## **4 *Summary and Recommendations***

TBD