



Revised plans for SRM: v2.2

- Critical features for WLCG
- Results of the May 22-23 workshop at FNAL
<https://srm.fnal.gov/twiki/bin/view/WorkshopsAndConferences/GridStorageInterfacesWorkshop>
 - SRM v2.2 definition geared to WLCG usage, but still compatible with other implementations
 - Some notions backported from SRM v3, others added for WLCG
- WLCG SRM “MoU”
<https://srm.fnal.gov/twiki/pub/WorkshopsAndConferences/GridStorageInterfacesWSAgenda/SRMLCG-MoU-day2.doc>
 - Needs some updates and polishing
- Schedule for implementation and testing
<https://srm.fnal.gov/twiki/pub/WorkshopsAndConferences/GridStorageInterfacesWSAgenda/Schedule.pdf>
- Measures to monitor progress and deal with issues



Critical features for WLCG

- Result of WLCG Baseline Services Working Group
 - <http://cern.ch/lcg/PEB/BS>
- Originally planned to be available by WLCG Service Challenge 4
 - Relevant implementations currently are incomplete and incompatible
- Features from version 1.1 + critical subset of version 2.1
 - File types
 - Space reservation
 - Permission functions
 - Directory functions
 - Data transfer control functions
 - Relative paths
 - Query supported protocols
- Details in presentation for LCG Internal Review of Services, June 9
 - <http://agenda.cern.ch/fullAgenda.php?ida=a062385>



SRM v2.2 MoU for WLCG

- Summarize agreed client usage and server behavior for the SRM v2.2 implementations used by WLCG applications
 - Servers can ignore non-WLCG use cases for the time being
- Clients
 - FTS, GFAL, lcg-utils
 - No direct usage of SRM by experiments
- Servers
 - CASTOR, dCache, DPM

- Stick with SRM v3 terminology for now, but with a WLCG understanding
 - TRetentionPolicy {REPLICA, CUSTODIAL}
 - OUTPUT is not used
 - TAccessLatency {ONLINE, NEARLINE}
 - OFFLINE is not used

- Tape-resident with system-managed disk cache
 - Tape1Disk0 == CUSTODIAL + NEARLINE
- Tape-resident with guaranteed copy on disk
 - Tape1Disk1 == CUSTODIAL + ONLINE
- Disk-resident, user-managed
 - Tape0Disk1 == REPLICA + ONLINE

- All WLCG files (SURLs) are permanent
 - Files can only be removed by the user

- Only deals with disk
 - Cache in front of tape back-end, and disk without tape back-end
 - Tape space considered infinite

- TapeNDiskM storage classes only require static reservations by VO admins
 - Can be arranged out of band without using the SRM interface (CASTOR)
 - Agreement between VO admin and SE admin will be needed anyway
 - Networks of main clients can be indicated (dCache)

- Dynamic reservations by ordinary users not needed in the short term
 - At least CMS want this feature in the medium term

- userSpaceTokenDescription attaches meaning to opaque space token
 - “LHCbESD” etc.

- srmPrepareToPut
 - To store a file in the space (i.e. storage class) indicated
 - WLCG clients will supply the space token
 - WLCG files are immutable, cannot be overwritten

- srmPrepareToGet
 - To prepare a file for “immediate” transfer or access
 - Recall from tape and/or copy to pool accessible by the client should now be done through srmBringOnline
 - WLCG usage excludes changing space or retention attributes of the file

- srmBringOnline
 - To indicate that a prepareToGet for the files is expected in the near future
 - A delay parameter can be used for further optimization
 - A prepareToGet could tie up resources, e.g. I/O movers in dCache
 - Signature very similar to that of prepareToGet
 - No TURLs are returned



srmChangeSpaceForFiles

- To change the storage class of the given files
 - Tape1Disk0 \leftrightarrow Tape1Disk1 (add/remove disk copy)
 - Tape0Disk1 \leftrightarrow Tape1DiskN (add/remove tape copy)
- To be decided which transitions shall be supported
- The SURL shall not be changed
 - Absolute path may change if SURL only contains relative path (as desired)
- Not required in the short term

- To get all metadata attributes for individual files, but only some for directories
 - Directory listings quickly become very expensive
- Directory listing use case would be to check consistency with file catalog
 - An implementation-dependent upper limit will apply for the time being
 - Use of the offset and count parameters requires further discussion



Schedule (1/3)

- WSDL and SRM v2.2 spec - June 6
 - Various inconsistencies have been fixed since
 - Discussion about the need for some unexpected changes w.r.t. v2.1
 - Stable version since June 19
- srmPrepareToGet, srmPrepareToPut at the same level of functionality as it is present now - June 20
 - Not technically challenging
 - Need 3 endpoints by the end of this period
 - dCache and DPM ready on June 26, CASTOR needs one or two days
 - Need a test suite – servers come with their own Java, C and C++ clients
 - FNAL srmcp - Apache Axis + Globus CoG Kit
 - Castor C++ client – gSoap + GSI plugin
 - DPM C client – gSoap + GSI plugin

- Compatibility 1 week after that - June 27
 - Alex Sim (LBNL) and Jiri Mencak (RAL) will run their test suites
- dCache srmCopy compatibility with DPM and CASTOR
srmPrepareTo(Get/Put) - work by Fermilab - July 4
- Space Reservation prerelease implementations - Sept 1
 - To coincide with SRM v2/v3 workshop at CERN, Aug 30 – Sept 1
- Space Reservation / Storage Classes - Sept 30 (optimistic)
 - Proper SRM or out-of-band way to reserve space
 - srmGetSpaceTokens
 - Modifications to srmPrepareToPut and srmCopy; srmPrepateToGet optional
 - srmRm, srmReleaseFiles (srmPurgeFromSpace not needed)
- Space Reservation may only work for special deployment configurations
 - Need to determine (per VO) if disk pools should be externally reachable



Schedule (3/3)

- srmBringOnline - Oct 6
- srmLs - return of space tokens is not required for October
- WLCG clients should follow the same schedule
 - Ready to be used as testers by the end of Sept
 - Will have several SRM test suites
 - Functionality, stress tests, error handling and resilience to (accidentally) “malicious” clients
- Integration week at RAL - Oct 9-13
 - Firm dates to be decided as milestone (by end of June)
- It could all work sufficiently by Nov 1
 - To allow v2.2 to become the standard SRM service (v1.1 for legacy apps)
 - Development of less urgent features will continue



Monitoring the progress

- Regular Friday phone conferences, mailing list
 - Involve developers, experiments, some sites
- Regular reports to GDB and WLCG MB
- Monthly report to overview committee
 - Line management of developers at CERN, DESY, FNAL, RAL
- Workshop Aug 30 – Sep 1, integration week Oct 9 - 13
- GDB will set up a group of site experts
 - Understand implications of storage classes and space reservation per VO
 - Disk pool layout
 - Networks