



# SRM status and plans

- What is SRM v2.2 ?
- Status of implementations
- Open issues
- Schedule



# What is the SRM?

- Client-server interface for Storage Resource Management
  - Spaces, files, directories, access, transfer, ...
  - See EGEE User Forum (2006/03/02) presentation for details
    - “Use of the Storage Resource Manager Interface”
    - <http://indico.cern.ch/sessionDisplay.py?sessionId=13&slotId=0&confId=286>
- Version 1.1 in widespread use
  - But implementations have subtle incompatibilities due to ambiguities in the “standard”
  - Various basic functionalities not defined
- Version 2.1 implementations incomplete or incompatible
- Version 3 definition not stabilized yet



# Critical features for WLCG

- Result of WLCG Baseline Services Working Group
  - <http://cern.ch/lcg/PEB/BS>
- 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
- Originally planned to be available by WLCG Service Challenge 4 (April)
  - Relevant implementations were incomplete or incompatible
  - Use cases have evolved since
- 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

- Mandated by WLCG Management Board after Mumbai (CHEP) workshop
- Designed during May 22-23 workshop at FNAL
  - Some notions backported from SRM v3, others added for WLCG
- Summarizes 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 LHC experiments
- Servers
  - CASTOR, dCache, DPM
  - Berkeley SRM
  - StoRM



# Storage classes

- Stick with SRM v3 terminology, 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



# Storage class usage

- A storage class instance is implemented as an SRM space
  - Created through srmReserveSpace or out-of-band by SE admin
  - These spaces are static, i.e. set up in advance by VO admin
    - Typically group-writable
    - Dynamic user spaces will come later
  - An SE usually will have multiple instances of any supported storage class
    - Different spaces for different VOs
    - Different spaces for different groups or different data types within a VO
- Clients will refer to spaces through space tokens or user descriptions
  - Explicit API and CLI options will be available
    - Else the default space for the VO will be used
  - To be supplied to srmPrepareToPut and srmCopy only
  - srmPrepareToGet does not need it in minimal WLCG model
    - A file can only be in a single space at a time



# SRM v2.2 / v3 meeting report

- SRM v2.2 meeting Aug. 29
  - <http://agenda.cern.ch/fullAgenda.php?ida=a063257>
  - Status reports and discussions
- SRM v3 meeting Aug. 30 – Sep. 1
  - <http://agenda.cern.ch/fullAgenda.php?ida=a062357>
  - Many discussions also very relevant to v2.2 (v2.3)
    - v3 functionality observed to be moving toward v2.2
- Details in WLCG Management Board presentation of Sep. 12
  - <http://agenda.cern.ch/fullAgenda.php?ida=a063265>



# Status of servers

- CASTOR
  - No significant development remaining for minimal WLCG model
- dCache
  - Space reservation alpha version ready
  - Directory and data transfer functions implemented
  - srmBringOnline in October
- DPM
  - All methods for minimal WLCG model essentially implemented
  - Pending issues, e.g. DB schema upgrade procedure
- Berkeley SRM
  - Almost all methods for minimal WLCG model implemented
- StoRM
  - All methods for minimal WLCG model expected around end September





# Unresolved issues

- Space ownership and ACLs vs. VOMS proxy contents
  - Not needed in the short term
  - Only production managers must be distinguished from ordinary users
    - Each implementation can do that based on DN for the time being
    - DPM could already take agreed VOMS Role into account
  - v3/v2.3 use subspaces for dynamic reservation by users
- ACLs on files and directories
  - Similar issues and timescales
- srmLs output format
  - Paths instead of SURLs
- File lifetimes
  - DB schema could already foresee v3/v2.3 enhancements



# Client status

- FTS
  - Use space token or description, if supplied, on Put and Copy
  - srmLs only on single files
  - bringOnline for pre-staging early next year
  - ACL propagation (also for leading directory creation) to be decided
  - Best effort support for v1.1
  - WLCG v2.2 support to be in release end of October
- GFAL/lcg-utils
  - Space token to be provided in explicit call to set the context
    - Not as extra arguments to the POSIX-like interface
  - ACL propagation to be decided
  - WLCG v2.2 support to be available end of October



# Berkeley SRM-Tester status

- Test client independent of any server implementation
- Basic functionality tests
- Advanced tests provoke failed requests
  - Check the server's subsequent behavior
- Basic and advanced interoperability tests with srmCopy
- Results published on web pages
  - <http://sdm.lbl.gov/srm-tester/>
- Stress tests implementation dependent on extra funding
  - Work part of GGF-GIN project



# Storage Classes working group

- Study deployment configurations of storage classes
  - Per SRM back-end
  - How to assign pools to the various storage classes needed by a VO
  - How to configure (subsets of) pools for LAN or WAN access
  - Try to devise common configurations for VOs, per site
- Tape1Disk1 implementation is not evident
  - The other two are standard
- Allowed transitions between classes are not evident
  - Want to avoid making physical copies
  - Explicit stageout Tape0 → Tape1 awkward to implement
    - Disallow?
- Space reservation refers to the pool from which the data is used
  - Data typically accessed from LAN only
  - WAN transfer pool takes up some HW, but is not exposed



# xrootd, quotas, httpg

- xrootd
  - BaBar have important use case for file update  $\leftrightarrow$  immutable files
    - File can change as long as not in catalog
    - Want prepareToPut with delay option to allow for updates etc.
  - Tested with modified Berkeley SRM
- Quotas
  - Still desired for v3/v2.3
  - Complementary to spaces
- httpg vs. https
  - httpg always delegates proxy, while only needed for srmCopy
    - Separate delegation service reduces security overhead
  - https is standard
    - Can use alternative implementations, less code to maintain



# GLUE schema changes

- New schema for SE to reflect properties of SRM v2.2 and higher
  - Storage Area → space
  - Access Policy info per SA per VO
    - Foresees quotas
  - Enhances Access Protocol info
  - Removes obsolete/unused notions
- Submitted to GLUE schema working group



# Plan of work

- Settle open issues in phone conferences + mailing lists
  - ACL semantics, ...
  - Avoid non-trivial changes to WSDL/spec for the time being
  - Collaborate with Storage Classes WG
- Continuous interoperability testing
  - Berkeley SRM-Tester
  - Test suite of Jiri Mencak (RAL) continued by Flavia Donno (CERN)
- Deploy v2.2 services on a few major sites by Nov. 1
  - Allow for large-scale stress tests by the VOs
  - A Service Challenge week in Jan./Feb. would be desirable
- Have v2.2 deployed on all relevant WLCG sites by spring
  - Keep v1.1 service for legacy applications
  - Define plan for enhancements beyond minimal WLCG model