SRM 2.2: tests and site deployment

30th January 2007

Flavia Donno, Maarten Litmaath
IT/GD, CERN

Tests executed

- S2 test suite testing availability of endpoints, basic functionality, use cases and boundary conditions, interoperability, exhaustive and stress tests.
  - **Availability**: Ping and full put cycle (putting and retrieving a file)
  - **Basic**: basic functionality checking only return codes and passing all basic input parameters
  - **Usecases**: testing boundary conditions, exceptions, real use cases extracted from the middleware clients and experiment applications.
  - **Interoperability**: servers acting as clients, cross copy operations
  - **Exhaustive**: Checking for long strings, strange characters in input arguments, missing mandatory or optional arguments. Output parsed.
  - **Stress**: Parallel tests for stressing the systems, multiple requests, concurrent colliding requests, space exhaustion, etc.
- S2 tests cron job running 6 times a day (overnight for US sites)

  http://cern.ch/grid-deployment/flavia
- In parallel, manual tests from GFAL/lcg-utils,FTS, DPM test suite and test suite executed at LBNL daily.

Tests executed

Basic MoU SRM methods

Needed now only for dCache!

MoU SRM methods needed by the end of 2007. Expected by the end of summer

Summary of S2 SRM v2.2 basic test - Saturday 27 January 2007 09:00am CET

SRM function	CERN CASTOR	FNAL DCACHE	CERN DPM	LBNL DRM	STORM
WLCG MoU SRM v2.2 methods					
Ping	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PrepareToPut	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
StatusOfPutRequest	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutDone	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PrepareToGet	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
StatusOfGetRequest	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
BringOnline	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
StatusOfBringOnlineRequest	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
AbortRequest	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
AbortFiles	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ReseaseFiles	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
GetRequestSummary	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
GetRequestTokens	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
GetTransferProtocols	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Ls	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Mkdir	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Rmdir	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Rm	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Mv	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ReserveSpace	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
StatusOfReserveSpaceRequest	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ReleaseSpace	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
GetSpaceTokens	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
GetSpaceMetaData	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ExtendFileLifeTime	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
WLCG MoU SRM v2.2 methods needed by end of 2007					
Copy	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
StatusOfCopyRequest	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ChangeSpaceForFiles	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log

Tests executed

Some stability. Some methods still not implemented

WSDL update

LHCC review announced

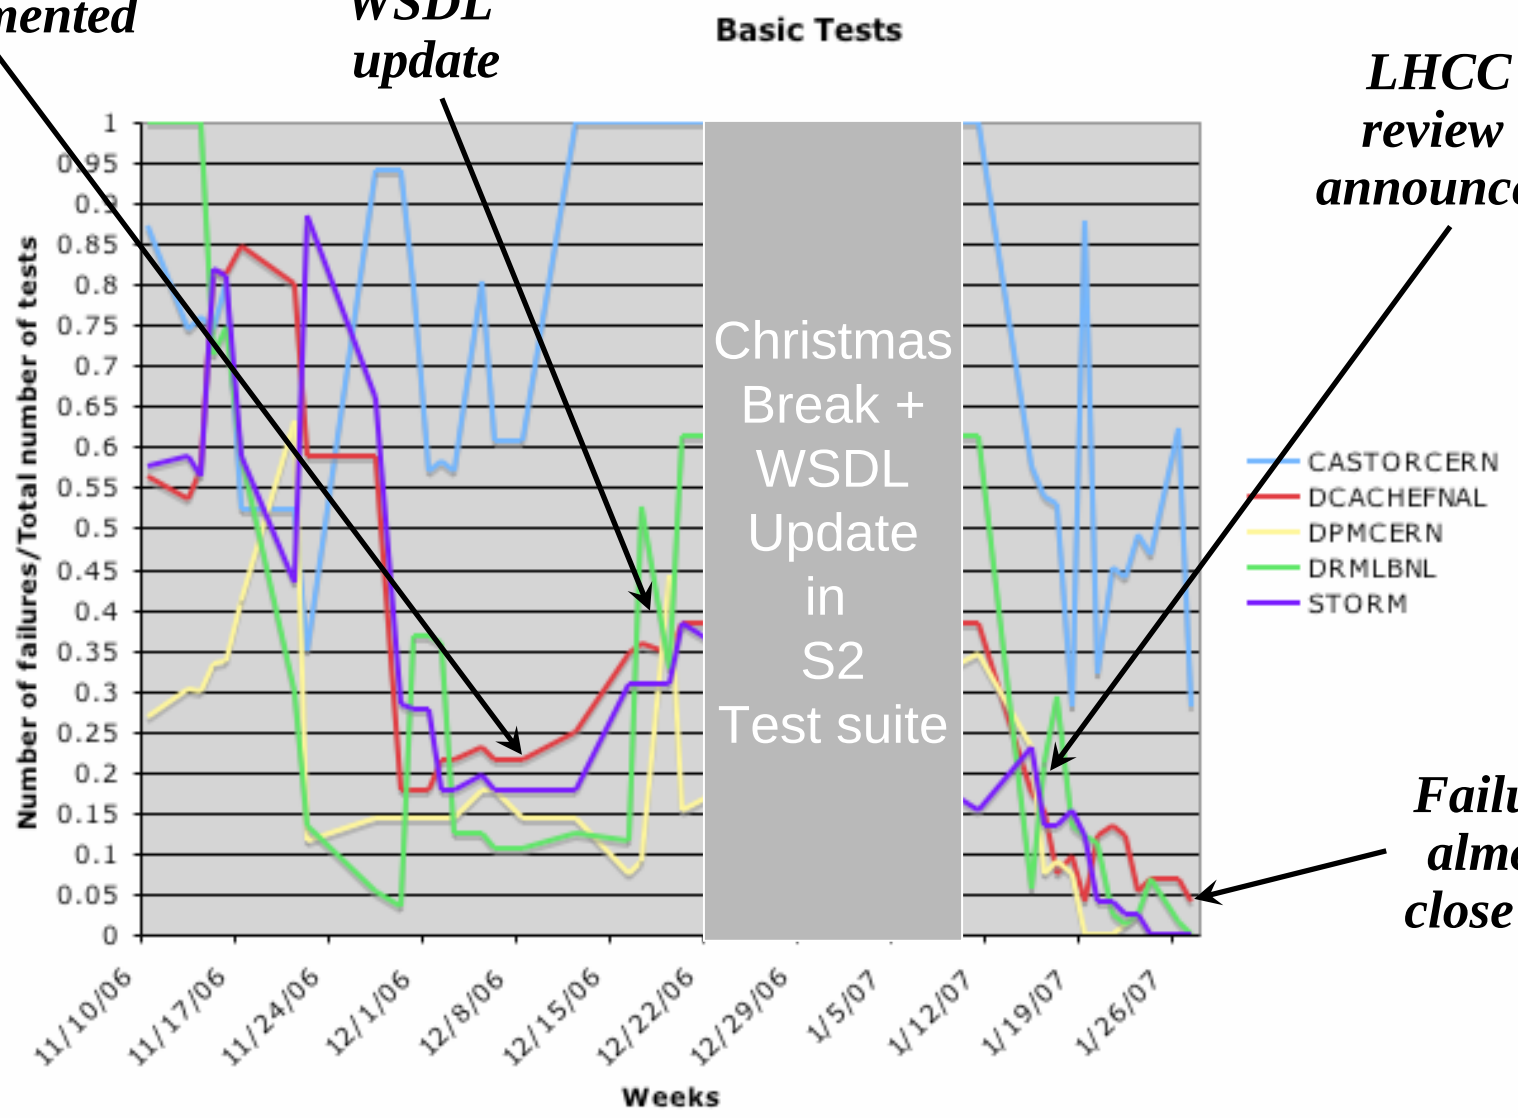

Failures almost close to 0

Tests executed

Availability

Summary of S2 SRM v2.2 availability test - Saturday 27 January 2007 09:00am CET

CERN CASTOR	FNAL DCACHE	CERN DPM	LBNL DRM	STORM
UP	UP	UP	UP	UP

Interoperability/ Cross Copy

Summary of S2 SRM v2.2 cross test - Saturday 27 January 2007 09:00am CET

SRM function	CERN CASTOR	FNAL DCACHE	CERN DPM	LBNL DRM	STORM
Copy Tests in PUSH mode					
CopyToFNALDCACHE	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyToCERNDPM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyToLBNLDRM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyToSTORM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Copy Tests in PULL mode					
CopyFromFNALDCACHE	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyFromCERNDPM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyFromLBNLDRM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
CopyFromSTORM	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log

UseCase

Summary of S2 SRM v2.2 use-case test - Saturday 27 January 2007 09:00am CET

SRM test	CERN CASTOR	FNAL DCACHE	CERN DPM	LBNL DRM	STORM
GetRemoved01	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
LsFullDetail	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
LsNonExistent	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Pin01	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Pin02	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Pin03	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Pin04	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PrepareToPutNEVER	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutGet01	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutNearline	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutNoOverwrite	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutOverwrite	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
PutRemoved01	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ReleaseFiles	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
ReserveSpace	StdOut	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Rmdir01	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Space00	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Space01	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log
Space02	StdOut	StdOut	StdOut Log	StdOut Log	StdOut Log
TokenSensCase	StdOut Log	StdOut Log	StdOut Log	StdOut Log	StdOut Log

Tests executed: status of the implementations

- DPM has been rather stable and in good state for the entire period of testing. Few issues were found (few testing conditions caused crashes in the server) and fixed. All MoU methods implemented beside Copy (not needed at this stage).
- DRM and StoRM: good interaction. At the moment all MoU methods implemented (Copy in PULL mode not available in StoRM). Implementations rather stable. Some communication issues with DRM need investigation.
- dCache: very good improvements in the basic tests in the last weeks. Implementation is rather stable. All MoU methods have been implemented (including Copy that is absolutely needed for dCache) but ExtendFileLifeTime. Timur has promised to implement it as soon as he gets back to US.
- CASTOR: The implementation has been rather unstable. The problems have been identified in transferring the requests to the back-end server. Therefore, it has been difficult to fully test with the basic test suite the implementation of the SRM interface. A major effort is taking place to fix these problems.

Plan

- Plan for 1Q of 2007 :
  - Phase 1: From 16 Dec 2006 until end of January 2007:
    - Availability and Basic tests
    - Collect and analyze results, update page with status of endpoints: https://twiki.cern.ch/twiki/bin/view/SRMDev/ImplementationsProblems
    - Plot results per implementation: number of failures/number of tests executed for all SRM MoU methods.
    - Report results to WLCG MB.
  - Phase 2: From beginning until end of February 2007:
    - Perform tests on use-cases (GFAL/lcg-utils/FTS/experiment specific), boundary conditions and open issues in the spec that have been agreed on.
    - Plot results as for phase 1 and report to WLCG MB.
  - Phase 3: From 1 March until "satisfaction/end of March 2007" :
    - Add more SRM 2.2 endpoints (some T1s ?)
    - Stress testing
    - Plot results as for phase 2 and report to WLCG MB.
- This plan has been discussed during the WLCG workshop. The developers have agreed to work on this as a matter of priority.

Grid Storage System Deployment (GSSD)

- Working group launched by the GDB to coordinate SRM 2.2 deployment for Tier-1s and Tier-2s
  - https://twiki.cern.ch/twiki/bin/view/LCG/GSSD
  - Mailing list: storage-class-wg@cern.ch
- Mandate:
  - Establishing a migration plan from SRM v1 to SRM v2 so that the experiments can access the same data from the 2 endpoints transparently.
  - Coordinating with sites, experiments, and developers the deployment of the various 2.2 SRM implementations and the corresponding Storage Classes.
  - Coordinating the Glue Schema v1.3 deployment for the Storage Element making sure that the information published are correct.
  - Ensuring transparency of data access and the functionalities required by the experiments (see Baseline Service report).
- People involved: developers (SRM and clients), providers, site admins, experiments

Deployment plan

- (A.) Collecting requirements from the experiments: more details with respect to what is described in TDR. The idea is to understand how to specifically configure a Tier-1 or a Tier-2: storage classes (quality of storage), disk caches (how many, how big and with which access), storage transition patterns, etc.
  - Now
- (B.) Understand current Tier-1 setup: requirements ?
  - Now
- (C.) Getting hints from developers: manual/guidelines ?
  - Now
- (D.) Selecting production sites as guinea pigs and start testing with experiments.
  - Beginning of March 2007 - July 2007
- (E.) Assisting experiments and sites during tests (monitoring tools, guidelines in case of failures, cleaning-up, etc.). Define mini SC milestones
  - March - July 2007
- (F.) Accommodate new needs, not initially foreseen, if necessary.
- (G.) Have production SRM 2.2 fully functional (MoU) by **September 2007**

Deployment plan in practice

- Detailed requirements already collected from LHCb, ATLAS, CMS. Still some reiteration is needed for CMS and ATLAS.
- Focus on dCache deployment first. Target sites: GridKA, Lyon, SARA. Experiments: ATLAS and LHCb.
  - Variety of implementations. Many issues covered by this exercise
- Compiling specific guidelines targeting an experiment: ATLAS and LHCb.
- Guidelines reviewed by developers and sites. Covering possibly also some Tier-2 sites.
- Working then with some of the T2s deploying DPM. Repeat the exercise done for d-Cache. Possible parallel activity.
- Start working with CASTOR if ready: CNAF, RAL.
- Define mini SC milestones with the experiments and test in coordination with FTS and lcg-utils developers.

Deployment plan in practice: the client perspective

- Site admins have to run both SRM v1.1 and SRM v2.2
  - Restrictions: SRM v1 and v2 endpoints must run on the same "host" and use the same file "path". Endpoints may run on different ports.
- Experiments will be able to access both old and new data through both the new SRM v2.2 endpoints and the old SRM v1.1
  - This is guaranteed when using higher level middleware such as GFAL, lcg-utils, FTS, LFC client tools. The endpoints conversion is performed automatically via the information system. New GLUE schema is needed (see examples in next slide).
- SRM type is retrieved from the information system
  - In case 2 versions found for the same endpoint SRM v2.2 is chosen only if space token (storage quality) specified. Otherwise SRM v1.1 is the default
  - FTS can be configured per channel on the version to use; policies can also be specified ("always use SRM 2.2", "use SRM 2.2 if space token specified",…)
- It is the task of the GSSD Working Group to define and coordinate configuration details for mixed mode operations.

===>>> It is possible and foreseen to run in mixed mode with SRM v1.1 and SRM v2.2, until SRM v2.2 is proven stable for all implementations.

Conclusions

- Much ***clearer*** description of SRM ***specifications***. All ambiguous behaviors made explicit. A few issues left out for SRM v3 since they do not affect the SRM MoU.
- Well ***established*** and agreed ***methodology*** to check the status of the implementations. Boundary conditions, use cases from the upper layer middleware and experiment applications will be the focus of next month's work. Monthly reports and problems escalation to the WLCG MB.
- A ***clear plan*** has been put in place in order to converge. Developers will work on it as a matter of priority
- Working with ***sites and experiments*** for the deployment of the SRM 2.2 and ***Storage Classes***. Specific guidelines for Tier-1 and Tier-2 sites are being compiled.
- It is not unreasonable to expect ***SRM 2.2 in production by September 2007***.
- The ***migration and backup plans*** foresee the use of a mixed environment SRM v1 and v2, where the upper layer middleware takes care of hiding the details from the users.