



Report from GSSD Storage Workshop

Flavia Donno
CERN

WLCG GDB
4 July 2007



The Workshop

- About 50 people attending the workshop: developers, site managers and experiment representatives
- The program focused on SRM development, testing and deployment strategies and plans
- In general, people felt that the workshop has provided good input and has contributed to progress
- Main outcome is that we established a detailed plan for SRM deployment in WLCG



The SRM Service

- Jamie Shiers presented the WLCG service schedule.
 - We discussed the implications of the LHC machine schedule and the experiment's Dress Rehearsals on the deployment of SRM v2.2 services
 - The conclusion was that we have 2 windows for SRM deployment: post CSA07 and early 2008.
- The Storage Service developers presented the status and plans for their services.
 - It was felt that we should now concentrate on what is really needed for SRM 2.2 deployment and stability of the services.
 - We must NOT dedicate time on providing new features.
- OSG presented its plan for testing and support of SRM v2 for both dCache (UCSD will test it) and BeStMan (no real customers at the moment). They do not intend to deploy Chimera at least in the short term.



The SRM Service

- S2 tests:
 - All implementations pass functionality and interoperability tests.
 - Stress tests are started with promising results (for details please check the slides in the agenda)
 - It was agreed that stress tests should continue till at least October 2007 on all development endpoint
- Discussions on VOMS-based ACLs have started and a draft document on functionalities needed has been circulated.
 - Normal non-VOMS proxy will continue to be supported
 - VOMS-based ACLs will be supported on files and directories
 - It was agreed that POSIX ACLs will be supported
 - Per VO a group/role will designate people with authority to change ACLs
 - After drafting a first proposal for ACLs, GSSD will get in contact with other groups in WLCG discussing the same matter (e.g. job priority wg) to get further input and understand implications.
 - Discussions on the matter will resume after CHEP (No coding required for the moment for implementations not providing the requested features)
- Next year (April ?), we will start organizing meetings to discuss what we learned from experience and draft a spec for a “WLCG SRM” v2.3.



The sites

- Sites reported on their experience installing, configuring and maintaining an SRM v2.2 service
 - dCache: various problems reported during installation and configuration by FZK and GridPP. Developers will address the problems reported in the next days. These issues will be collected in one of the GSSD web pages together with the solutions provided, as it is done now for the SRM issues.
 - DPM: Small inconsistency reported with dpm-updatespace. A bug in Savannah has been > open. The fix is available in the next DPM release
 - StoRM: no clear documentation. More information needed for sizing the system installing different components on different machines.
 - CASTOR: some experience only at RAL and CNAF.



The sites

- CNAF reported on performance tests comparing dcap, GPFS, rfio, xrootd
 - GPFS performs well and satisfies CNAF Tier-1 requirements.
 - They performed tests using StoRM. They demonstrated a throughput of 360MB/sec for several hours.



The experiments

- The experiments presented their plans and requirements for testing SRM v2.2 in pre-production
 - They all stressed that they have no resources to dedicate to the job and the easiest way to test SRM v2.2 is to include it in their “normal production chain”.
 - Therefore, they all expressed the wish to access production data from the SRM v2.2 endpoints made available for the tests.
 - Grid production computing resources should be used during this exercise.



The experiments

- Other comments:
 - In general, sites do not have extra resources to dedicate to pre-production exercises. Extra resources should therefore come from the current production resources in use or such resources should be considered in the MoU.
 - Given the limited human and hardware resources available and testing time window, the experiments accept to introduce SRM v2.2 in production whenever an “acceptable confidence level” has been reached and make it work there, before the LHC starts.



The roll-out plan: a summary

- Following recent progress with SRM v2.2 implementations, we are now able to define a roll-out schedule, together with acceptance criteria and intermediate milestones
- The experiments' and LHC schedule give two slots for such deployment:
 1. October 15-31 for a small number of key sites that can be supported with priority (and other sites "best-effort")
 2. Early January 2008 for remaining sites
- High-level tools will be modified to set v2.2 as the default version of SRM. This will happen by October 15th, 2007
- The status will be followed on a regular basis by the WLCG MB
- It is understood that full scale production shake-down will only occur once the main sites have upgraded and that this will take several months. This means that the schedule should not be delayed if the services are to be ready well ahead of 2008 data taking



The criteria

- Production releases of the SRM v2.2 implementations are required, together with documentation and an agreed support model (including training)
 - Documentation sufficient for “intermediate” site to setup and configure a working system
- These versions should pass the existing functionality and stress tests sustained for a period of one week
 - Failures permitted on a maximum of one day; no more than 15% performance degradation over the week
- The experiment tests (next) should also be passed
 - Definition: end-July
 - Execution: end-August
- Follow-up: September 1st DM sessions at WLCG workshop



Experiment tests

Experiment	To be tested with priority	Technical contact	MB contact	Comments
ALICE	TBD	TBD	Federico Carminati	
ATLAS	FTS 2.0 tests from CERN to BNL, FZK, IN2P3		Dario Barberis	Performance 50% greater than current
CMS	FTS 2.0 transfers to 2-3 sites		Matthias Kasemann	On channels that are currently not production quality. PhEDEx quality plots
LHCb	FTS 2.0 transfers and data access to 2-3 sites		Nick Brook	Requires production CPU & storage resources at these sites



The plan and resources

- The detailed roll-out plan will be published by the end of this week on the GSSD twiki page
- We are gathering information about available hardware and human resources at sites participating in the tests
- Specific information about configuration requirements for Tier-1s and Tier-2s will be published as well

PLEASE, STAY POSTED!

<https://twiki.cern.ch/twiki/bin/view/LCG/GSSD>