

SRM 2.2: an update



2th July 2007

Flavia Donno
CERN, IT/GD



Tests

- All implementations pass basic tests
 - <https://twiki.cern.ch/twiki/bin/view/SRMDev>
- Use-case test family enhanced with even more tests:
 - **CASTOR**: Passes all use-cases.
 - Disk1 implemented by switching off garbage collector (not gracefully handled by CASTOR). Proper Disk1 support will come in the next Castor version (2.1.4, to be finished in July, deployed the next month(s)).
 - PutDone slow. *Fixed*
 - **dCache** : Passes all use-cases except 2 (srmGetRequestTokens, srmGetSpaceMetadata) *Fixed*
 - In Tape1Disk0 allocated space decreases when files are migrated to tape *Fixed*
 - **DPM**: Passes all use-cases except 1 (missing space garbage collector) *Fixed*
 - Garbage collector for expired space available with next release of DPM (1.6.5 in certification). *Fixed*
 - **StoRM**: Passes all use-cases except 1 *Fixed*
 - It is possible to use expired space *Fixed*
 - **BeStMan (not used in WLCG)**: Several use-cases not passed *Fixed*
 - Tokens are case insensitive *Fixed*
 - Some calls not compliant to the specs as defined during the WLCG Workshop in January 2007. *Fixed*



Tests

- **Stress tests** started on all development endpoints using 9 client machines. Small server instances are preferred in order to reach easily the limits.
- **First Goals:**
 - Understand the limits of the instance under test
 - Make sure it does not crash or hang under heavy load
 - Make sure that the response time does not degrade to an “unreasonable” level
- **Further goals:**
 - Make sure there are no hidden race-conditions for the calls that are the most used
 - Understand server tuning
 - Learn from stress testing
- In Parallel, stress-testing activities are on-going by the EIS team with GSSD input

First Results



- **CASTOR:**
 - Race conditions found. Working with developers to address problems.
 - Good handling of heavy-load: requests are dropped if server busy (the client can retry)
 - Response time for the requests being processed is good.
- **dCache:**
 - Authorization module crash
 - Server very slow or unresponsive (restart cures the problem)
 - Working with developers to address problems
- **DPM:**
 - No failures
 - Good handling of heavy-load: requests are dropped if server busy (the client can retry)
 - Response time for the requests being processed is good.
- **StoRM**
 - Response time degrades with load. However, it recovers after the crisis.
 - Working with developers to address problems
- **BeStMan**
 - Server unresponsive under heavy load. It does not resume operations when load decreases.
 - Working with the developers to address problems



Deployment

- 17 endpoints in the S2 testbed (<https://twiki.cern.ch/twiki/bin/view/LCG/GSSDendpoints>)
- Focusing on key Tier-0/1 sites (CERN, CNAF, BNL, FZK, IN2P3)
 - These sites are gaining experience on SRM v2.2 configuration and management.
 - They will move SRM v2.2 in production once the confidence level has been reached and experiments have performed tests successfully.
 - GSSD is coordinating this effort.
- Endpoints (test instances) available in pre-production **for the experiments to test.**
 - First tests executed moving data from CASTOR@CERN (SRMv1) production to BNL (dCache), CNAF (StoRM), CERN (DPM) (all SRMv2) using FTS 2.0: ATLAS and LHCb - OK!
 - On-going: Tests full chain transfer+data access for LHCb involving part of the LHCb Tier-1 sites: CERN, CNAF, FZK (CASTOR and dCache instances)
 - Some agreements for tests with CMS involving FTS 2.0 and Tier-1 SRM v2.2 sites
 - Successful tests of data access using SRM v2.2 and high-level utilities: GFAL, lcg-utils. Some bugs found and fixed.
- **Testing by experiments is crucial to declare SRM v2.2 ready for production.**



Conclusions

- The basic functionality, use case and interoperability tests are complete.
- It is essential to deploy SRM v2.2 at Tier-1s by the end of this year to ensure smooth operations at LHC start-up.
- We (GSSD) are preparing a plan that targets a few sites at a time and allows for coexistence of SRM v1 and SRM v2 for some time.
- This is integrated with the experiment dress-rehearsals, which provide a good opportunity for the validation of SRM v2.2.
- BUT - In general, Tier-1 sites do not have resources (both human and hardware) to ensure parallel deployment of SRM v1 and v2
- AND - Experiments do not have human resources to dedicate to the validation of SRM v2.2 in production (interference with current dress-rehearsal preparation activities).
- The schedule for SRM v2.2 validation and deployment in production is extremely tight and will only be achieved if developers, sites and experiments give priority to an agreed roll-out plan.
- Tier-2s can upgrade and switch to SRM v2.2 without major problems, although some central coordination is required.