

SRM 2.2: an update



3th July 2007

Flavia Donno
CERN, IT/GD



Tests

- All implementations pass functionality tests
 - <https://twiki.cern.ch/twiki/bin/view/SRMDev>
- Use-case test family enhanced with more tests.
- Remaining issues:
 - CASTOR:
 - 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)).
 - dCache/StoRM :
 - Very different ways to "reserve" space for the experiments and provide space token descriptions. VO Managers or site administrators ? Need a clear solution.



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



Outcome of the Storage Workshop

- Need to prove that SRM v1 is "better" than SRM v2. The metrics are:
 - New functionalities (functionality tests)
 - Stability
 - Stability will be proven stress testing endpoints for a considerable amount of time (one week) and measuring the server response
- Major commitment from sites and developers to roll-out SRM v2.2 in production by the end of 2007.
- A clear roll-out plan has been agreed.
- It was noted that sites do not have resources for PPS services. They should be taken out from MoU agreed resources for a VO.



SRM v2.2 roll-out plan

Plan for Tier-1s:

- Focus on few sites: BNL, FZK, IN2P3, NDGF, SARA, RAL, CNAF, CERN. Start with dCache sites.
- Differentiate between a testing phase (till October 15th, 2007) and a deployment phase (from 15 Oct. 2007 to January 2008)
- Have dCache test instances properly configured for ATLAS and LHCb for middle July 2007. All confirmed availability but BNL. Probably also SARA is in question.
- Have DESY properly configured with dCache by middle of next week. Start sustained stress tests on these 2 instance, using several certs at the same time and performing a mixture of SRM v2.2 requests. Prove stability over a period of at least a week under heavy load. Test to be performed over the summer.
- Test SRMv1-v2 interoperability for all possible implementations with high level tools/apis using entries in production catalogues with multiple certs. Check as well accessibility/manageability of SRM v2 data through SRM v1 endpoints. Test to be performed over the summer.



SRM v2.2 roll-out plan

Plan for Tier-1s:

- Push experiment active testing by next week, July 9th 2007. Start with whatever site is available. (ATLAS with FZK ?).
- Continue stress tests on development instances and experiment testing till middle October 2007.
- Roll-out patches as they come out (a strategy for rolling out patches is being established by the developers). GSSD will coordinate installation of patches at key sites.
- Start deployment of SRM v2.2 at FZK, IN2P3, DESY October 15th, 2007 (in agreement with the experiments). The other key Tier-1 dCache sites will follow.



SRM v2.2 roll-out plan

Plan for Tier-1s:

- CASTOR at CERN will make available a test endpoint for LHCb testing. This can be properly configured and ready by middle of July 2007.
- CASTOR will be deployed at RAL for ATLAS by middle of July 2007. RAL is planning to support LHCb by August 2007. They can also setup SRM v2 endpoints to production instances.
- CASTOR at CNAF will also make available a test instance for LHCb/ATLAS by middle of July 2007.
- Stress and experiment tests as for dCache. S2 stress only development instances.
- Acquire experience and start wide deployment of CASTOR in production by October 2007.



SRM v2.2 roll-out plan

Plan for Tier-2s:

- Tier-2s using DPM can migrate to SRM v2 as of now. The configuration can be coordinated centrally by GSSD with the input from the experiments.
- Tier-2s using dCache can upgrade as soon as a confidence level for configuration, management, and operations has been reached at dCache key Tier-1 sites.
- Tier-2s using StoRM can migrate to SRM v2 as soon as a space token description can be associated to Storage Area (middle July 2007).
- Roll-out patches as they come out with the coordination of GSSD.
- Deployment of SRM 2.2 should be completed by end of January 2008.



SRM v2.2 roll-out plan

Plan for high-level tools:

- By the 15th of October new releases of higher level tools (GFAL, lcg-utils) will offer the possibility of configuring the version of SRM to use by default. FTS 2.0 offers already the possibility to configure the version of SRM to use per channel.



SRM v2.2 roll-out plan

Coordination of support:

- Building experience at Tier-1 sites to offer support to Tier-2s
- Exercise GGUS channels already during the testing phase
- The developers will be the last level of the escalation process during normal operations but not during the testing and startup phase.
- Organize SRM specific tutorials/workshops covering configuration and management for all Storage implementations. This will happen after October 2007.