

SRR implementation and deployment

Julia Andreeva

WLCG Operations Coordination meeting

03.09.2020

What is SRR, a short reminder

- SRR stands for Storage Resource Reporting proposal developed as a follow up of the WLCG accounting review at the (pre-)GDB in April 2016. [More info](#)
- Several rounds of ‘consensus building’ with experiments, storage providers and sites. Still evolving.
- Prototyping phase has been agreed at the GDB in October 2017. [More info](#)
- [The latest specification of the SRR format](#)

Two important dimensions of SRR

- Storage topology description (SRR for topology)
 - Protocols
 - Storage shares equivalent/similar to SRM space quotas (path with quota defined, dedicated for a particular usage, accounted separately, not overlapping in terms of space)
- Non-SRM Storage Space Accounting enabled (SRR for accounting)
 - Non-SRM protocol to query used/free space
 - JSON file which includes both topology description and accounting data accessible through the http or xrootd protocols with the credentials of a particular experiment

SRR for accounting

- Two possible solutions are foreseen:
 - Enable at least one non-srm protocol (http, xrootd) to query used/free space for every storage share/quota included in the storage topology description file
 - Extend storage topology description file with accounting data (used/free space and time stamp of the measurement)
- Can be used both for operations (request on demand) and storage space accounting system (will collect data hourly). Foreseen frequency of update order of 30 minutes.
- Already used by ATLAS in production (via RUCIO) and the WLCG Storage Space Accounting (WSSA).

SRR readiness for different storage implementations

SRR for DPM (~49 sites)

- Is ready and deployed at the most DPM sites. [More info](#)
- SRR generation and publishing is implemented as a native DPM functionality. No additional actions are required by the site administrators

SRR for dCache (~ 43 sites)

- Has been implemented , but after deployment to the dCache sites a major problem with empty storage shares section at many sites has been discovered.
- The fix has been recently prototyped by the dCache team and has been deployed at DESY. After validation of the provided solution dCache would need to be upgrade at all sites where SRR problem has been detected
- The collecting and providing of the storage accounting information is a native dCache functionality. However publishing of this information is performed by the external script which should be run as a cron job

SRR for StoRM (~22 sites)

- Is ready and deployed at the most DPM sites. [More info](#)
- SRR report is generated by a native storm component, the storm-info-provider. Publishing of the file is performed by the an external script. No further development in this respect is planned

SRR for EOS

- Is ready and deployed at CERN. There are some issues which we are following up with EOS experts. However, these issues are rather related to a complex configuration of storage space for CERN EOS. Which areas to be accounted, which are not, what is pledged what is not, what to do with overcommitment, etc...
- SRR generation is a native EOS functionality. File is generated in a standard location and can be fetched with xrootd protocol. So in principle, no actions are required by the site admins.
- When CERN EOS issues are sorted out will continue to deploy it at other EOS sites

SRR for ECHO

- Deployed at RAL
- SRR generation is implemented with a cron which runs the “ceph df --format=json” command and puts it into the SRR format.

SRR for xrootd

- First prototype has been created for xRootd
- Did not progress much so far since this solution most probably would be required for the US sites for which we do not have topology information in CRIC. Therefore we hope to get help from US ATLAS and US CMS

SRR for US ATLAS and US CMS sites

- Solutions mentioned earlier are mostly used at the European sites.
- There is an ongoing effort in order to understand what we should do with the US T2 sites
- Many of the US CMS T2 sites run xrootd on top of HDFS storage. Discuss with CMS how we can address this configuration. One of the possible solutions is to enable xrootd space queries in the xrootd-hdfs plugin.
- Will schedule a meeting with US ATLAS next week

How we make sure that SRR is valid

- We are working on SRR checker which will run every hour or so and check whether SRR is uptodate , whether is is accessible and whether its content is correct. The SRR checker will update CRIC with SRR status and version of the storage implementation
- First version should be ready and deployed by the end of September