

Storage Space Accounting introduction

The goal of the WLCG Storage Space accounting project is to enable the high level overview of the total and available space provided by the WLCG infrastructure. What is required is a possibility to account total used and total free space for all distinct storage areas (equivalent to the space quotas in SRM) available to the experiments.

Work on storage space accounting is going in several directions:

- Enable description of the storage topology and all storage areas which have to be accounted separately
- Enable possibility to query storage space accounting information for all kind of storage implementation
- Enable the data flow of the accounting information from the information source to the data repository and setting up user interface and APIs for data retrieval

The storage space accounting work is done in close collaboration with the WLCG Data Management Steering group. The WLCG Data Management Steering group will coordinate with the storage providers in order to enable storage resource reporting. The WLCG accounting is just one of the use cases where storage resource reporting is required. The latest version of the storage resource reporting proposal can be found here [?](#).

WSSA service is available here [?](#).

More details:

- Jan 25, 2018 status update [?](#)
- Jan 16, 2019 GDB presentation [?](#)

Validation of data provided by WSSA

- January 2019 comparison WSSA vs ATLAS DDM Dashboard
- January 2019 comparison WSSA vs EGI accounting prototype

Description of the storage topology and all storage areas which have to be accounted separately

Storage systems should provide total used and total free space for all distinct space quotas available to the experiments. Therefore, description of all those distinct space quotas or storage areas which have to be accounted separately is required. In difference with the accounting information, this is pretty static information. The new WLCG topology and configuration system CRIC is foreseen as a place where this information will be stored and exposed via UI/API. The final format of the storage topology description and how this information will be provided by the sites hosting the storage service is still under discussion and will be followed up in the collaboration with the Data Management Steering Group. More details can be found in the document [?](#).

The latest version of the json format specification can be found here

The previous version in the googledoc [?](#)

Query storage space accounting information for all kind of storage implementation

The minimum requirement is to have two numbers free and used space for every space quota/storage area. The accessibility of these numbers depends upon the storage system type, the protocol, and configuration decisions relating spaces quotas with the namespace. In reality, there are three possibilities, gridFTP, HTTP and xrootd. A storage system should implement resource reporting in at least one protocol. The relevant numbers will be made available through the gfal2 interface.

Alternative solution for providing accounting information is to provide it in the json file, similar to the one which describes storage topology , but extended with the accounting data.

More details can be found in the latest version of the document [↗](#).

Implementation of the data flow for storage space accounting information

The central WLCG Storage Space Accounting (WSSA) service every 30min/1hour does the following:

- Queries the topology source, currently experiment-specific sources are used. Cric will be used in the future
- Queries distributed storage instances in order to retrieve accounting info per every space quota/storage area
- An alternative way would be to retrieve both topology and accounting data from the json file described below

This information is then published to the central repository (ES backend) using standard MONIT infrastructure. Grafana is used for the UI implementation.

Implementation of this data flow has been already prototypes using available information sources. For ATLAS and LHCb - SRM queries, for ALICE - ALICE ldapsearch queries. In their turn ATLAS and LHCb are using SRM for their internal storage space accounting, while ALICE uses xrootd queries. There are couple of known issues with xrootd space queries. For some storage configuration they can return too high numbers, double counting, which is in most cases fixed on the ALICE side . That is why it was decided to use ALISE ldap instead of raw xrootd queries. Another issye is with dcache storage for which currently xrootd space queries do not work. The problem is being followed up by the dCache developers.

SRR implementation by the storage middleware providers

SRR implementation is being followed up with all storage middleware providers. Mailing list for people involved in this activity is srr-implementation@cernNOSPAMPLEASE.ch

SRR implementation status

Storage type	Implementation status	Deployment status	Storage contact
DPM	Implemented starting from version 1.10.3. Requires DPM re-configuration with DOME enabled	DPM upgrade Task Force	Fabrizio Furano & Oliver Keeble
dCache	First prototype is ready ↗ for testing and validation. Path for storage shares is missing, requires a bit more work	dCache upgrade Task Force	Paul Millar
EOS	Enabled at CERN. Some issues to be fixed. Not yet ready for deployment to other sites	CERN only	Andreas Peters
STORM	First prototype at CNAF is ready for testing and validation, STOR-1139 ↗ , documentation ↗	-	Andrea Ceccanti
xRootd	First prototype provided beginning of November	-	Wei Yang

Examples of implementation

- DPM example for Prague
- EOS example for CERN
- STORM example for CNAF
- dCache example for Desy
- Xrootd example/template for SLAC
- Xrootd example/template for OU

Meetings with storage middleware providers

- Meeting on storage topology description 24.09.2018 [↗](#)
- Meeting SRR specification 06.11.2018 [↗](#)

-- JuliaAndreeva - 2018-10-04

- xrootd.srr.slac.json: Xrootd SRR json at SLAC
- OU-SRR-example.txt: OU-SRR-example.txt
- srr-json.py.txt: srr-json.py.txt
- quota_file_srr: quota_file_srr

This topic: LCG > StorageSpaceAccounting
Topic revision: r17 - 2021-03-17 - HorstSeverini



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback